

A rank-based refinement of ordinal efficiency and a new (but familiar) class of ordinal assignment mechanisms

Clayton R. Featherstone*

This version: October 21, 2011

JOB MARKET PAPER

PRELIMINARY AND INCOMPLETE

Abstract

A feasible assignment is *rank efficient* if its distribution over agents' ranks of their allocations cannot be stochastically dominated by that of another feasible assignment. Rank efficiency implies ordinal efficiency (and hence ex post efficiency); however, the converse is not true, even if we are free to give different weights to each agent when tabulating the rank distribution. A class of simple linear programming mechanisms always yields rank efficient assignments and can generate any rank efficient assignment. These mechanisms have been previously seen in the field and are currently used by Teach for America to assign teachers. Using data from a match at Harvard Business School, Featherstone and Roth (2011) empirically show that under truth-telling, rank efficient mechanisms can significantly outperform random serial dictatorship and the probabilistic serial mechanism. Although rank efficiency and strategy-proofness are theoretically incompatible, we show that rank efficient mechanisms can admit truth-telling as an equilibrium in low information environments. Rank efficient mechanisms are also strongly related to the competitive equilibrium mechanisms of Hylland and Zeckhauser (1979) which we are able to leverage to show that envy-freeness and rank efficiency are theoretically incompatible.

*Harvard Business School

JEL: D00, D02, D04, D51

Keywords: one-sided matching, ordinal efficiency, linear programming, general equilibrium, assignment market

1 Introduction

In an ordinal assignment problem, the goal is to design an effective mechanism that maps agents' submitted ordinal preferences into an assignment that matches each agent with one object. Such mechanisms are widely used in settings as diverse as school choice, housing allocation, machine scheduling, and military posting. To gauge success, it is common practice to look to how many people were assigned objects that they ranked highly. For instance, when school districts announce the results of a choice-based student placement system, they report things like how many students were assigned to one of their top three choices and how many were unmatched (NYC Department of Education 2009, San Francisco Unified School District 2011). Motivated by this near ubiquitous phenomenon, we consider evaluating the efficiency of ordinal assignment mechanisms by looking at the market-wide distribution of agents' ranks for their allocations. Remaining agnostic about how to compare different ranks, we call a feasible assignment *rank efficient* if its rank distribution cannot be first-order stochastically dominated by that of another feasible assignment.

Prior to this paper, two efficiency concepts dominated the literature. Ex post efficiency simply means that once a deterministic assignment is implemented, no agents would agree to trade or would want to claim unassigned objects. Random serial dictatorship (drawing an ordering from some fixed distribution and letting the agents choose their objects in that order) is ex post efficient and strategy-proof (Zhou 1990, Abdulkadiroğlu and Sönmez 1998). It is also commonplace in the field. It is not, however, efficient from the interim perspective in which the mechanism has chosen a lottery over assignments, but has yet to determine which of these assignments is to be implemented. From such a perspective, it is convenient to think about agents possessing a bundle of tradable probability shares in the different objects. Interim inefficiency then means that there are mutually profitable opportunities for agents to trade shares. The contribution of Bogomolnaia and Moulin (2001) was to identify this problem and to suggest a refinement of ex post efficiency that deals with it, which they called ordinal efficiency. They then introduced the probabilistic serial mechanism and showed that it is always ordinally efficient. Finally, they proved that strategy-proofness and ordinal efficiency are theoretically incompatible.³ Unfortunately, while ordinal efficiency is an appealing concept, the probabilistic serial mechanism is never

³Actually, they prove that there is no mechanism that is strategy-proof, ordinal efficient, and gives the same expected allocation to all agents who submit the same preference. The last condition rules out serial dictatorships, which are ordinally efficient by virtue of making the definition vacuous: agents can't trade across *states* of the world when there is only one *state*.

seen in the field. This presents an institutional puzzle, since good mechanisms often evolve on their own (Roth 1984, 1991, Abdulkadiroğlu and Sönmez 2003*b*).

As Bogomolnaia and Moulin (2001) presented an interim refinement of ex post efficiency, this paper presents an ex ante refinement of ordinal efficiency. Before going further, however, it makes sense to give some intuition for why such a refinement is needed. Consider a five agent, five position example, where the circles and boxes represent the deterministic assignments \boxed{x} and $\odot(x)$.

$$\begin{aligned}
 1 &: \boxed{a} \succ \odot(e) \\
 2 &: \odot(a) \succ \boxed{b} \\
 3 &: \odot(b) \succ \boxed{c} \\
 4 &: \odot(c) \succ \boxed{d} \\
 5 &: \odot(d) \succ \boxed{e}
 \end{aligned}$$

Both of these assignments are ex post efficient, and in fact, they are also ordinally efficient.⁴ Even so, \boxed{x} gives one agent her first choice and four agents their second choice, while $\odot(x)$ has four first choices and only one second choice. Rank efficiency formalizes the intuition that, if we don't have any reason to differentiate between one agent getting his first choice or another, then $\odot(x)$ is a better assignment.

Although this sort of interpersonal comparison may give pause to the theorist, it is hard to imagine a policy-maker that wouldn't prefer $\odot(x)$ to \boxed{x} . Indeed, when faced with an assignment whose rank distribution is worse than expected, policy-makers often want to “correct” assignments by looking for non-Pareto-improving rearrangements in which, by some definition, the good outweighs the bad. For instance, when Teach for America⁵ assigns its teachers to schools in different regions of the country, they start with a computerized match, and then spend a solid week manually looking for trade cycles that help teachers who have been assigned to low ranked regions, often at the expense of moving other teachers out of their top choice. After Harvard Business School (HBS) used random serial dictatorship to assign first year MBAs to overseas programs, administrators were concerned about the number of students who were assigned to an unpopular country even though they had ranked it last.⁶ HBS

⁴In Section 5.3 we will show that, on the domain of deterministic assignments, ordinal efficiency is equivalent to ex post efficiency.

⁵Teach for America is a nationwide non-profit that sends mostly new college graduates to teach in at-risk public schools. The author is able to speak to Teach for America's matching methods because he and Al Roth are actively involved with helping the organization to streamline its assignment process. See Section 5.4 for more.

⁶This paper discusses the HBS match in Section 6. Again, the author is able to speak to what

wanted to look for swaps where one MBA who somewhat liked the unpopular country, but wasn't assigned there, could trade places with another MBA who strongly disliked the unpopular country, but was assigned there. Only after being reminded that this would undermine the promise of strategy-proofness made to the students did administrators reluctantly accept the original assignment.⁷

Returning to the institutional puzzle mentioned earlier, one could interpret the fact that the probabilistic serial mechanism is absent from the field as an indication that institutional evolution has selected strategy-proofness over efficiency. The two anecdotes of the previous paragraph indicate otherwise, and in fact, rank efficiency can be characterized in terms of non-Pareto improving trade cycles. Consider scoring a trade cycle by giving values to the ranks and looking at the change in rank-value of all of its agents. For instance, if an agent in the cycle starts with his first choice and ends up with his third choice, then he contributes $v_3 - v_1 < 0$ points to the valuation of the cycle. An assignment is rank efficient if and only if there is some vector v such that all possible trade cycles yield a negative value. In other words, the sort of "correction" processes mentioned in the previous paragraph are not hard to rationalize as a search for rank efficiency. The characterization of the rank efficient mechanisms is built on a similar principle. Consider scoring every assignment as follows: for every agent (in expectation) that gets his 1st choice, the assignment gets v_1 points, for every agent (in expectation) that gets his 2nd choice, the assignment gets v_2 points, and so on. The mechanism that picks the feasible assignment with the highest score is a rank efficient mechanism. Such a process is easily automated with a simple linear program, and in fact, linear programming mechanisms, although not as prevalent as random serial dictatorships, can be found in the field. Roth (1991) mentions that the mechanisms used to match house-officers⁸ to hospitals in Cambridge and London are linear programs, and, since spring of 2011, Teach for America has also used a linear program.⁹

Thus far, rank efficiency has been described as a positive concept; however, it can also be motivated as a normative concept. If the v vector mentioned in the previous

happened behind the scenes because he and Al Roth are actively involved with the design of the match.

⁷Indeed, strategy-proofness is largely about making it safe for agents to reveal that they somewhat like objects that most other agents strongly dislike.

⁸Roughly speaking, house-officers are a U.K. analog to medical residents in the American market (Roth 1984).

⁹One might object that the author and Roth suggested this mechanism; however, in Section 5.4, we argue that Teach for America adopted a linear programming assignment method, not because it was something new, but because it duplicated the non-automated process of the past in much less time.

paragraph is interpreted as an *assumption* about the cardinal preferences of the agents and about how utility should be compared across agents, then the linear programming mechanisms we described coincide with maximizing expected utility from behind the veil of ignorance. In this sense, rank efficiency is an ex ante refinement of ordinal efficiency.

Institutional evolution seems to have framed the conflict between strategy-proofness and efficiency as a decision between random serial dictatorship and linear programming. As a market designer, how does one choose which mechanism is the right one? Do rank efficient mechanisms ever make sense, or are they “mistakes” made by uninformed policy-makers? Two exercises inform these concerns. First, in low information environments, à la Roth and Rothblum (1999), we show that rank efficient mechanisms admit a truth-telling equilibrium. This result lines up well with conversations we had with Teach for America about the wedge between strategy-proofness and efficiency. When asked why they weren’t worried about manipulation, their response was that teachers only applied once, were geographically separated, and knew little about what regions are popular and how the mechanism is run. The second exercise that informs whether it makes sense to consider rank efficient mechanisms uses empirical data from the HBS match. Under truth-telling, the rank efficient mechanism can yield more than a 15% improvement in the number of MBAs who can be assigned to a first or second choice country (out of 11) when compared to the assignment given by random serial dictatorship (Featherstone and Roth 2011). Although these gains might be undermined by strategic preference manipulation, the exercise shows that in a worst-case scenario, the costs of strategy-proofness can be quite large. Together, these exercises indicate that in certain environments, it would be a mistake to not at least consider using a rank efficient mechanism.

Finally, the paper considers a class of ordinal mechanisms based on Hylland and Zeckhauser (1979). These mechanisms assume cardinal preferences to rationalize the submitted ordinal preferences, give each agent a budget of fiat¹⁰ money, and then calculate a competitive equilibrium assignment. Such mechanisms always yield ordinally efficient assignments, and in fact, they yield a refinement of ordinal efficiency that generalizes rank efficiency by allowing different agents to be weighted differently when the rank distribution is tabulated. Hence there is a mapping between generalized rank efficient mechanisms¹¹ and competitive equilibrium mechanisms. Perhaps

¹⁰That is, the money only exists within the mechanism and is not intrinsically valued by the agents. Its only worth is to purchased probability shares within the mechanism.

¹¹As we just discussed, the rank efficient mechanism gives an assignment v_k points for assigning any agent his k^{th} choice. The generalized rank efficient mechanism weights agents according to some

surprisingly, a competitive equilibrium where all agents have the same budget does not correspond to an assignment supported by the rank efficient mechanism that places equal weights on each agent. In this sense, procedural fairness means different things under the two types of mechanisms. Equal budgets in a competitive equilibrium mechanism corresponds to justice based on envy-freeness, (Varian 1974, Dworkin 1981) while equal agent weights in a rank efficient mechanism corresponds to justice based on a version of the Rawlsian veil of ignorance (Harsanyi 1975).¹²

The rest of the paper is essentially organized into three parts. In the first part (Sections 2-5), we start by introducing the model, ex post efficiency and ordinal efficiency, and then demonstrate that rank efficiency is a refinement of both. Then, we go on to characterize rank efficient mechanisms, interpret what they are doing, and argue that they are used in the field by telling the story of how Teach for America assigns its teachers. In the second part (Sections 6-7) we show that rank efficient mechanisms are worth thinking about by demonstrating two things. First, by looking at data from the HBS overseas match, we show that under truth-telling, rank efficient mechanisms could yield big efficiency gains. Then, we show that, although strategy-proofness and rank efficiency are theoretically incompatible, rank efficient mechanisms can admit a truth-telling equilibrium in low information environments. In the third part (Sections 8-9) we generalize rank efficiency and show that the generalization is closely related to the competitive equilibrium mechanisms of Hylland and Zeckhauser (1979). Leveraging this result, we then interpret the theoretical incompatibility of rank efficiency and envy-freeness as a wedge between two important concepts of justice. Finally, we conclude by discussing how the rank efficient mechanisms should enter into the discussion about the costs of strategy-proofness.

2 The model

Consider assigning each agent a from set \mathcal{A} to exactly one object o from set \mathcal{O} . Further, let there be q_o copies of each object. Sometimes the set of objects will include a special “null” object, \emptyset , that denotes an agent’s outside option. We model \emptyset as a good that is never scarce, that is, $q_\emptyset = |\mathcal{A}|$. When agents have no outside

vector $(\alpha)_a$: the value of assigning an agent a to his k^{th} choice is worth $\alpha_a \cdot v_k$.

¹²Note that we are maximizing *expected utility* from behind the veil. If we were to follow Rawls (1972), then we would maximize a maximin objective which would lead to a concept of justice that more closely resembles that of Dworkin.

option,¹³ we don't include the null object in \mathcal{O} . Whether there is a null object or not, we will require that there are enough objects that every agent can be feasibly matched, that is $\sum_{o \in \mathcal{O}} q_o \geq |\mathcal{A}|$.¹⁴

A **deterministic assignment** is a function that maps agents to objects **feasibly**, that is, each agent is only mapped to one object, and no more than q_o agents are assigned to any given $o \in \mathcal{O}$. We can represent such an assignment as an $|\mathcal{A}| \times |\mathcal{O}|$ matrix x where $x_{ao} \in \{0, 1\}$, $\sum_o x_{ao} = 1$, and $\sum_a x_{ao} \leq q_o$, for all $a \in \mathcal{A}$ and $o \in \mathcal{O}$. To be clear, $x_{ao} = 1$ means that agent a is assigned to object o ; $x_{ao} = 0$ means that agent a is not. A **random assignment** is a lottery over deterministic assignments, which can be represented as the corresponding convex combination over deterministic assignment matrices. As such, random assignment matrices will have a similar structure to deterministic assignment matrices, except $x_{ao} \in [0, 1]$. By the extension of the Birkhoff-von Neumann theorem (Birkhoff 1946) put forth in Budish et al. (2011), we can assert that any such matrix represents some lottery over deterministic assignments.¹⁵ We call this a **lottery representation** of the random assignment matrix, and the deterministic assignments the representation's **support**. For the rest of the paper, we will frequently use the freedom afforded to us by the Budish et al. theorem to focus on matrix representations.

Moving to the individual agent, we refer to the a^{th} row of a random assignment x , x_a , as agent a 's **allocation**. Each agent a is endowed with an ordinal preference \succsim_a over \mathcal{O} ; note that indifferences are allowed. As the framework of this paper is laid out, however, we will often have occasion to express agents' preferences in terms of rank functions. An agent a 's rank function, $r_a(\cdot)$, is a mapping from \mathcal{O} to $\{1, \dots, |\mathcal{O}| + 1\}$.¹⁶ When preferences are strict, there are several equivalent ways to define the standard rank function: $r_a(o) = |\{o' \in \mathcal{O} | o' \succsim_a o\}|$, or $r_a(o) =$

¹³The obvious example is military postings. Perhaps a more subtle example is public school assignment. In San Francisco, students are not required to rank all schools, but if they cannot be assigned to a school they ranked, they are generally given an administrative assignment. In this sort of situation, failing to rank all schools is equivalent to the student saying to the school district, "beyond what I ranked, you can fill out the rest of my rank-order list for me."

¹⁴More generally, if some agents are not allowed to match to certain objects, we just need that there exists some feasible deterministic assignment from agents to objects. If we consider a bipartite network where links represent feasible allocations, then a necessary and sufficient condition for the existence of such an assignment is that the network meets the condition of Hall's Theorem (Hall 1935), that is, for any subset of the agents, the cardinality of the union of the sets of objects to which they could feasibly match must be greater than or equal to the cardinality of the subset of agents in consideration.

¹⁵This lottery need not be unique. Generally, we don't have to worry about this, as all such lotteries induce the same lottery over objects for each agent, but sometimes this subtlety is important. See Remark 1.

¹⁶We include $|\mathcal{O}| + 1$ for technical reasons that will become clear in the next paragraph.

$|\{o' \in \mathcal{O} \mid o' \succ_a o\}| + 1$, or even $r_a(o) = |\{[o'] \in \mathcal{O} / \sim_a \mid o' \succeq_a o\}|$.¹⁷ With indifference, however, these definitions are no longer equivalent. Consider $\succeq_a = a \succ b \sim c \succ d$. The ranks of (a, b, c, d) under these definitions are, respectively, $(1, 3, 3, 4)$, $(1, 2, 2, 4)$, and $(1, 2, 2, 3)$. Another ambiguity in how to think of the rank function concerns the null object, \emptyset . One could of course treat \emptyset just like any other object; however, policy-makers often report the number of unassigned agents as a separate category that is, in some sense, worse than any other rank. When considering **individually rational** assignments, that is assignments where no agent is ever assigned to something less preferred than \emptyset , we can model this by setting $r_a(\emptyset) = |\mathcal{O}|$ and for all $o \prec_a \emptyset$, $r_a(o) = |\mathcal{O}| + 1$. In Section 7, we will see that the specifics of the mapping from \succeq_a to $r_a(\cdot)$ (which we call the **ranking scheme**) can affect incentives for truth-telling, but for now, the theory can be built around any of these definitions. We will only require that $o' \succ_a o \Leftrightarrow r_a(o') < r_a(o)$, $o' \sim_a o \Leftrightarrow r_a(o') = r_a(o)$.

Finally, define an **ordinal assignment mechanism** to be a mapping from submitted preferences to random assignments. In this paper, we focus on ordinal mechanisms because their simplicity and prevalence in the field.

3 From ex post to interim

We can think of the efficiency of assignments from several perspectives. Ex post efficiency makes sure that there are no Pareto improving rearrangements once a deterministic assignment is implemented. Interim efficiency takes a step back and ensures there are no Pareto improving rearrangements of object probability shares at the point where we have a random assignment, but have yet to choose which deterministic assignment in its support is to be realized. When cardinal preferences are known, both of these concepts are well defined and easy to implement.

With ordinal mechanisms, however, interim efficiency is a tricky concept, since we don't have enough information to determine agents' complete preferences over lotteries. The great innovation of Bogomolnaia and Moulin (2001) was to realize that there are some interim rearrangements that would be profitable regardless of the cardinal preferences that rationalize the reported ordinal preferences. If an assignment weakly improves the allocations of all agents in the first-order stochastic dominance sense (strictly for one agent), then regardless of the cardinal preferences, we have a Pareto improvement. Ordinal efficiency ensures that these sorts of improvements are

¹⁷ \mathcal{O} / \sim_a is the set of indifference classes of \mathcal{O} with respect to \sim_a ; hence, $|\{[o'] \in \mathcal{O} / \sim_a \mid o' \succeq_a o\}|$ is the number of indifference classes whose objects are weakly preferred to o .

taken advantage of.

For the rest of this section, we will briefly go over ex post and interim efficiency in the ordinal setting, with a view toward showing that rank efficiency, the topic of the present paper, dovetails nicely with previous literature. In Section 5, we will show that rank efficiency is a natural ex ante refinement of ordinal efficiency. Before proceeding, however, we briefly address indifferences. As mentioned in the previous section, rank efficient mechanisms will be capable of dealing with indifferences. This is also true of the ex post and ordinally efficient mechanisms we are about to introduce, but for the sake of simplicity and familiarity, we will introduce them in the context of strict preferences. References for the reader interested in understanding how indifferences can be incorporated are given, and new results concerning indifferences are relegated to the Appendix.

3.1 Ex post efficiency

Ex post efficiency means the random assignment can be represented by a lottery such that, regardless of which deterministic assignment is realized, there won't be a rearrangement of the objects that all agents weakly prefer (strictly for one). Formally,

Definition 1. A feasible deterministic assignment \tilde{x} is said to **ex post dominate** another deterministic assignment x if $\tilde{x}_a \succeq_a x_a$ for all $a \in \mathcal{A}$, and there is some a where the preference is strict. A feasible deterministic assignment is **ex post efficient** if it is not ex post dominated. A random assignment \tilde{x} is **ex post efficient** if the support of some lottery representation of the assignment consists entirely of ex post efficient deterministic assignments.

Remark 1. It is possible for an ex post efficient assignment to have a lottery representation whose support is not entirely ex post efficient. See Abdulkadiroğlu and Sönmez (2003a) for an example.

Remark 2. The definition of ex post efficiency for random assignments can be equivalently expressed in terms of a state-by-state domination concept. See the Section A in the Appendix for more.

Now we move on to characterize the class of ex post efficient mechanisms, the random serial dictatorships.

Definition 2. Serial dictatorship maps the reported strict preferences of the agents, (\succ_a) , and a permutation π of $(1, \dots, |\mathcal{A}|)$ to a deterministic assignment ac-

ording to to the recursion relations $\mathcal{O}_0 = \mathcal{O}$, $x_{a_{\pi(k)}} = \max_{\succ_{a_{\pi(k)}}} \mathcal{O}_{k-1}$, and $\mathcal{O}_k = \mathcal{O}_{k-1} \setminus x_{a_{\pi(k)}}$. Denote the assignment that results from these recursions as $SD[\succ; \pi]$.¹⁸

Definition 3. Random serial dictatorship maps the reported strict preferences of the agents, (\succ_a) , and a distribution over permutations of $(1, \dots, |\mathcal{A}|)$, $p(\pi)$, to the random assignment that results from a lottery that picks $SD[\succ; \pi]$ with probability $p(\pi)$. Denote the resulting random assignment as $RSD[\succ; p]$.

Proposition 1. *x is ex post efficient relative to strict preferences $\succ \Leftrightarrow \exists$ a distribution over permutations p such that $x = RSD[\succ; p]$*

Proof. Bogomolnaia and Moulin (2001) show that a deterministic assignment x is ex post efficient $\Leftrightarrow \exists \pi$ such that $x = SD[\succ; \pi]$.

(\Leftarrow) $RSD[\succ; p]$ is a lottery over ex post efficient assignments, by construction.

(\Rightarrow) The lottery representation of x whose support is ex post efficient gives us the distribution over orderings required. \square

Remark 3. Random serial dictatorships with indifferences have been discussed in the literature (Svensson 1994, 1999, Bogomolnaia, Deb and Ehlers 2005), albeit from more theoretical perspective. See Section B in the Appendix for more about how indifferences can be practically incorporated into random serial dictatorship.

So the random serial dictatorships always yield ex post efficient assignments, and any ex post efficient assignment can be supported by some random serial dictatorship. Random serial dictatorship is also strategy-proof.

Proposition 2 (Bogomolnaia and Moulin 2001). *If the distribution over orderings is fixed before agents submit their preferences, then random serial dictatorship is strategy-proof.*

Strategy-proofness and ex post efficiency are strong reasons to expect the random serial dictatorships to evolve naturally, and indeed, they are ubiquitous.

3.2 Ordinal efficiency

To understand why we might need a refinement of ex post efficiency, it helps to look at a situation in which ex post efficiency is too permissive. Consider the following example, adapted from Che and Kojima (2010).

¹⁸For deterministic assignments, we will sometimes abuse notation and let x_a denote both the a^{th} row of x and the object that x_a represents.

$$\begin{aligned}
A_1 &: a \succ b \succ c \\
A_2 &: a \succ b \succ c \\
B_1 &: b \succ a \succ c \\
B_2 &: b \succ a \succ c
\end{aligned}$$

There is only one copy of objects a and b , while there are two copies of object c . Random serial dictatorship relative to the uniform distribution over all orderings yields the random assignment¹⁹

$$\begin{array}{rcccc}
& & a & b & c \\
A_1 &: & 5/12 & 1/12 & 1/2 \\
A_2 &: & 5/12 & 1/12 & 1/2 \\
B_1 &: & 1/12 & 5/12 & 1/2 \\
B_2 &: & 1/12 & 5/12 & 1/2
\end{array}$$

which, by construction, must be ex post efficient. However, there is a mutually beneficial trade: A_1 and A_2 can trade their shares in b for the shares of a held by B_1 and B_2 .

$$\begin{array}{rcccc}
& & a & b & c \\
A_1 &: & 1/2 & 0 & 1/2 \\
A_2 &: & 1/2 & 0 & 1/2 \\
B_1 &: & 0 & 1/2 & 1/2 \\
B_2 &: & 0 & 1/2 & 1/2
\end{array}$$

Clearly, all agents prefer this random assignment, since they have traded probability shares of a less preferred object for shares in a more preferred object. The contribution of Bogomolnaia and Moulin (2001) was to demonstrate this problem and to suggest a natural ordinal extension of interim efficiency which remedies it.

Definition 4. A random assignment \tilde{x} **ordinally dominates** another assignment x

¹⁹To see this, note that only the first two agents in the ordering get a or b ; hence, every agent gets c half the time. For an agent a not to get his preferred object, he must be second in the ordering (1/4 of the time) and, conditional on that, the first agent in the ordering must take a 's first choice (1/3 of the time). $(1/4) \times (1/3) = 1/12$.

if for all agents a , \tilde{x}_a weakly stochastically dominates x_a with respect to \succsim_a , that is

$$\sum_{o \succsim_a o'} \tilde{x}_{ao} \geq \sum_{o \succsim_a o'} x_{ao}, \forall a \in \mathcal{A}, o' \in \mathcal{O}.$$

with strictness for at least one agent-object pair. An assignment is called **ordinally efficient** if there is no other assignment that ordinally dominates it.

The first thing to notice is that ordinal efficiency is a refinement of ex post efficiency.

Claim 1 (Bogomolnaia and Moulin 2001). x is ordinally efficient \Rightarrow x is ex post efficient, but x is ex post efficient $\not\Rightarrow$ x is ordinally efficient.

Proof. The leading example of this section shows the ($\not\Rightarrow$) part. For the (\Rightarrow) part, note that if x is a lottery whose support includes an ex post inefficient assignment, then replacing that assignment with its dominator would stochastically improve all agents from the interim perspective. \square

Ordinally efficiency also removes the subtlety of ex post efficiency mentioned in Remark 1; all lottery representations of ordinally efficient assignments have an ordinally efficient support.

Claim 2 (Abdulkadiroğlu and Sönmez 2003a). x is ordinally efficient \Rightarrow the support of any lottery representation of x must be entirely ex post efficient.

Proof. The same proof from the previous claim works. Note that the converse, however, is not true; see Abdulkadiroğlu and Sönmez (2003a) for a counterexample. \square

To produce ordinally efficient assignments, we look to the set of **simultaneous eating mechanisms**. Since they are somewhat complicated to explain precisely, and are not the focus of the present paper, we will limit ourselves to a rough description.²⁰ Essentially, over the “time” interval $[0, 1]$, each agent takes object shares from his most preferred remaining object at a pre-defined rate (given by his “eating speed” function²¹). If the object an agent is eating is exhausted, then he continues to eat at his next most preferred remaining object. The algorithm terminates when “time” gets to 1. Note that each agent is doing this *simultaneously*. By the Birkhoff-von Neumann theorem mentioned earlier, we can calculate a lottery over deterministic

²⁰For a more precise definition of this class of mechanisms, see Bogomolnaia and Moulin (2001).

²¹This function integrates over time to one so that an agent will consume a total of one share of probability on the interval $t \in [0, 1]$.

assignments that provides all agents with the lottery over objects that they consumed. It turns out that the simultaneous eating mechanisms are related to ordinally efficient assignments in the same way that random serial dictatorships are related to ex post assignments.

Proposition 3. *x is ordinally efficient relative to strict preferences $\succ \Leftrightarrow x$ can be generated by some simultaneous eating mechanism*

Remark 4. This theorem can be extended to a class of mechanisms that deal with indifferences (Katta and Sethuraman 2006); again, we assume strict preferences for expositional ease.

A commonly referenced simultaneous eating mechanism is the one in which all agents have identical eating speed functions, known as the **probabilistic serial** mechanism. For intuition, note that probabilistic serial would produce the ordinally efficient outcome we mentioned in the leading example of this section. Simultaneous eating mechanisms are not generically strategy-proof, that is, some agents can improve their allocation by deviating from truthful preference revelation. Finally, note that while the random serial dictatorships are quite common in the field, simultaneous eating mechanisms are not. This presents an institutional puzzle: if ordinal efficiency is theoretically appealing, then why don't we see ordinally efficient mechanisms?²² As we will show in the next section, the rank efficient mechanisms are ordinally efficient, which provide a potential resolution.

Our discussion in this section has introduced an example where there is a wedge between efficiency and strategy-proofness. Empirically, we might wonder what sort of efficiency gain the market designer purchases by forgoing strategy-proofness. This line of analysis is pursued further in Section 6.

4 Rank efficiency and the rank-value mechanisms

Intuitively, ordinal efficiency can sometimes be too permissive. Consider the following example where the boxes and circles represent potential deterministic assignments \boxed{x} and $\circledast x$.

²²Serial dictatorships (one fixed dictatorship ordering) are ordinally efficient, but vacuously so, as all ex post deterministic assignments are also ordinally efficient (see Section 5.3). More precisely then, the puzzle is why we don't ordinally efficient mechanisms that satisfy equal treatment of equals, that is, mechanisms that give the same allocation to all agents who submit the same preference.

$$\begin{array}{l}
1: \boxed{a} \succ \textcircled{e} \\
2: \textcircled{a} \succ \boxed{b} \\
3: \textcircled{b} \succ \boxed{c} \\
4: \textcircled{c} \succ \boxed{d} \\
5: \textcircled{d} \succ \boxed{e}
\end{array}$$

Both are ordinally efficient (and hence ex post efficient) deterministic assignments;²³ however, in some sense, \textcircled{x} seems better, as it gives three more agents their first choice and three fewer their second choice when compared to \boxed{x} . Interpersonal utility comparisons are, of course, tricky business, but given that we only have information about the agent's ordinal preferences, it is hard to imagine a policy-maker remaining agnostic about whether \textcircled{x} is better than \boxed{x} .

4.1 Rank efficiency

Consider the cumulative frequency distribution of ranks received by the agents in a market. Formally, define the **rank distribution of assignment x** to be

$$N^x(k) \equiv \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \mathbf{1}_{\{r_a(o) \leq k\}} \cdot x_{ao}$$

$N^x(k)$ is the expected number of agents who get their k^{th} choice or better under assignment x .

Definition 5. A random assignment x is **rank-dominated** by a feasible assignment \tilde{x} if the rank distribution of \tilde{x} first-order stochastically dominates that of x , that is, $N^{\tilde{x}}(k) \geq N^x(k)$ for all k (strict for some k). A feasible random assignment is called **rank efficient** if it is not rank-dominated by any other feasible assignment.

Although it is not immediately clear, rank efficiency is a refinement of ordinal efficiency. Formally,

Proposition 4. x is rank efficient \Rightarrow x is ordinally efficient, but x is ordinally efficient $\not\Rightarrow$ x is rank efficient.

Proof. The leading example of this section showed the $(\not\Rightarrow)$ part. For the (\Rightarrow) part, note that weakly first-order stochastically improving all agents (strict for one) will

²³In fact, on the domain of deterministic assignments, ordinal efficiency is equivalent to ex post efficiency. See Claim 9.

necessarily lead to a first-order stochastic improvement in the rank distribution, since the global rank distribution is a sum over the agents' personal rank distributions. \square

Also note that, in a similar way to ordinally efficient assignments, any lottery distribution of a rank efficient assignment must have a completely rank-efficient support.

Claim 3. The support of any lottery representation of any rank efficient random allocation must consist entirely of rank efficient deterministic allocations.²⁴

Proof. Say that the support of some lottery representation of an rank efficient assignment contained a rank dominated deterministic assignment. Then that assignment is rank dominated by the compound lottery that, in the support, replaces that rank dominated deterministic assignment with its dominator, a contradiction. \square

Schematically, thus far we have shown that

$$\begin{pmatrix} \text{ex post} \\ \text{efficiency} \end{pmatrix} \begin{matrix} \Leftarrow \\ \not\Rightarrow \end{matrix} \begin{pmatrix} \text{ordinal} \\ \text{efficiency} \end{pmatrix} \begin{matrix} \Leftarrow \\ \not\Rightarrow \end{matrix} \begin{pmatrix} \text{rank} \\ \text{efficiency} \end{pmatrix}$$

that is, rank efficiency is a refinement of ordinally efficiency, just as ordinal efficiency is a refinement of ex post efficiency. But why this particular refinement? In short, there is a good deal of evidence that policy-makers gauge their success by the rank distribution. For instance, school districts with ordinal choice mechanisms almost always report the distribution of ranks received as a measure of the quality of the match (NYC Department of Education 2009, San Francisco Unified School District 2011). This lends credence to rank efficiency as a positive concept that is at least implicitly considered in the field. In Section 5, we will further support this interpretation by showing that we see rank efficient mechanisms in the field. We will also give some support for rank efficiency as a normative concept, as it can be interpreted as a social welfare function, modulo a few assumptions.

First though, we return to a simplification of this section's leading example in order to concretely show how random serial dictatorship and probabilistic serial can fail to produce rank efficient assignments. Redefine \boxed{x} and \textcircled{x} relative to the example that follows. We also consider the output of the uniform random serial dictatorship and probabilistic serial mechanisms: their random assignments are listed in Table 1.

²⁴Note that on the domain of deterministic assignments, rank efficiency is not equivalent to ex post efficiency. See Claim 10.

x^{U-RSD}				x^{PS}					
	a	b	c	d		a	b	c	d
1:	1/2	0	1/2	0	1:	1/2	0	1/2	0
2:	1/2	1/6	0	1/3	2:	1/2	1/4	0	1/4
3:	0	5/6	1/6	0	3:	0	3/4	1/4	0

Table 1: Random assignments in the simplified example

	$N^x(1)$	$N^x(2)$	$N^x(3)$
\boxed{x}	1	3	3
$\odot x$	2	3	3
x^{PS}	$7/4 = 1.75$	$11/4 = 2.75$	3
x^{U-RSD}	$11/6 \approx 1.83$	$8/3 \approx 2.67$	3

Table 2: Rank distributions from the simplified example

$$\begin{aligned}
1: & \quad \boxed{a} \succ \odot c \succ d \\
2: & \quad \odot a \succ \boxed{b} \succ d \\
3: & \quad \odot b \succ \boxed{c} \succ d
\end{aligned}$$

The rank distributions of the four assignments are listed in Table 2. We can see that $\odot x$ rank dominates the other three assignments, since, column-by-column, it has a larger rank distribution. Of course, this will not always be the case (rank efficiency only guarantees that no other assignment rank dominates); in this example, it is due to the following claim, whose proof illustrates how to show that an assignment is the unique rank efficient assignment.

Claim 4. In the simplified example, $\odot x$ is the unique rank efficient assignment.²⁵

Proof. Claim 3 means that we can just show that $\odot x$ is the unique *deterministic* rank efficient assignment. First we show that the rank distribution of $\odot x$ dominates any other feasible rank distribution. The three agents' first choices only cover two objects; hence, $N^x(1) \leq 2$, and since there are only three agents, $N^x(2) \leq 3$. $\odot x$ hits

²⁵One might argue that this fact casts doubt on rank efficiency as a concept. What if we had good reason to care more about 1 than 2 and 3? In Section 8.1 we will introduce a generalized version of rank efficiency in which agents can be weighted differently when the rank distribution is tabulated. Everything we have shown thus far has an analog to a world in which the weighting over students is fixed, but not uniform.

this upper bound. Now we show that no other deterministic assignment attains the bound. Any assignment that attains the bound must give b to 3 and split a between 1 and 2. Since b was given to 3, if we give any a to 2, we will get $N^x(2) < 3$. Thus, (x) uniquely attains the bound, demonstrating that it is the unique rank efficient assignment. \square

Hence, a corollary of the previous claim is that any rank efficient mechanism will choose (x) . A natural next step is to look for a method of calculating the rank efficient assignments for a more general setting.

4.2 The rank-value mechanisms

We will first introduce the rank-value mechanisms, and then we will show that these mechanisms are to rank efficiency what the random serial dictatorships were to ex post efficiency. Begin by defining a **valuation** to be a sequence $(v_k)_{k=1}^{|\mathcal{O}|+1}$ of strictly positive real numbers such that $v_k > v_{k+1}$ for all $k \in \{1, \dots, |\mathcal{O}|\}$. v_k will be, in some sense, the “value” that the rank-value mechanism places on k^{th} choice allocations, so we have no need for more dimensions than $|\mathcal{O}|+1$ (which is the worst rank that anyone could possibly give to an object).²⁶ Also note that the strictness of the inequality is essential for many upcoming results; we must place a strictly higher value on k^{th} choice allocations than on $(k+1)^{\text{th}}$ choice allocations.

Definition 6. The **rank-value mechanism with respect to valuation v** (or the **v -rank-value mechanism** for short) maps agent rank orderings to a maximizer of the following linear program.

$$\begin{aligned} \max_x \quad & \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} v_{r_a(o)} \cdot x_{ao} \\ \text{s.t.} \quad & \sum_a x_{ao} \leq q_o, \quad \forall o \in \mathcal{O} \\ & \sum_o x_{ao} \leq 1, \quad \forall a \in \mathcal{A} \\ & x_{ao} \geq 0, \quad \forall o \in \mathcal{O} \\ & \quad \quad \quad \forall a \in \mathcal{A} \end{aligned}$$

We say that the assignments in the arg max are **supported by the v -rank-value mechanism**.

²⁶See the end of Section 2 for a reminder about why we might want an $(|\mathcal{O}|+1)^{\text{th}}$ rank.

A few straightforward properties of rank-value mechanisms follow almost immediately.

Claim 5. For any assignment x in the arg max, $\sum_o x_{ao} = 1$ for all $a \in \mathcal{A}$.

Proof. All elements of the valuation vector are strictly positive by definition, so all agents will be assigned a full unit of probability shares. \square

Having the constraint be an inequality instead of an equality will allow us to use duality theorems later on.

Claim 6. The v -rank-value mechanism is **individually rational**, that is if $\emptyset \in \mathcal{O}$, it never assigns an agent to something she likes less than \emptyset .

Proof. Say that there is an assignment in the arg max that assigns agent a to something he likes less than \emptyset . Then, moving a to \emptyset increases the objective and is feasible, since $v_k > v_{k+1}$ and $q_\emptyset = |\mathcal{A}|$. \square

Individual rationality allows us to ignore agent preferences below \emptyset . Now, we introduce a property that is useful for the characterization in Section 5.3. A mechanism is **non-wasteful** if it always chooses an assignment such that $x_{ao} > 0 \Rightarrow [o' \succ_a o \Rightarrow \sum_a x_{ao} = q_o]$. Non-wasteful assignments don't have unassigned objects that could be used to improve an agent's welfare.

Claim 7. The v -rank-value mechanism is non-wasteful.

Proof. It is feasible for an agent to swap any shares he has with unclaimed shares, and doing so increases the objective. \square

Finally, although we won't often need to know which maximizer of the linear program is chosen for the theory we are to develop, this selection is important for implementation. Consider the set of deterministic assignments. Since this set is finite and countable, we can rank its elements. Let the **tiebreaker function**, $\tau(x)$, denote the rank assigned to the deterministic allocation x . Our selection rule will then be to choose the deterministic assignment in the linear program's arg max with the lowest τ value. The following claim tells us that this is enough to yield a well-defined mechanism

Claim 8. The arg max of the v -rank-value mechanism is the convex hull of some set of deterministic assignments.

Proof. This is a linear optimization on a compact set, so the arg max is not empty. Consider a lottery representation of some assignment in the arg max. By the linearity of the problem, it must be that the deterministic assignments in the support are also in the arg max; otherwise, we could improve the objective by dropping them. Hence, any assignment in the arg max is a convex combination of deterministic assignments that are also in the arg max. Also by the linearity of the problem, a convex combination of any set of assignments in the arg max must also be in the arg max, since the convex combination will have the same objective value as its components. \square

So, the tiebreaker procedure will always choose some element of the arg max, and if we draw the tie-breaker ordering from some distribution, the tie-breaker procedure is effectively implementing a lottery representation of some random assignment in the arg max. Of course how we choose the tiebreaker can affect the incentives of the mechanism. In this paper, we will focus on tiebreaker functions that are drawn from a distribution and fixed before the agents submit their preferences. The most obvious of these procedures is to choose the tiebreaker uniformly at random from the set of all tiebreaker functions. This is the same as finding all deterministic assignments in the arg max and picking one uniformly at random, which in turn is the same as implementing the random assignment that is the centroid of the arg max. Of course, other more *ad hoc* methods seem likely to work, although we do not attempt to prove so.

Regardless of the tiebreaker, however, the family of rank-value mechanisms and the set of rank-efficient assignments are very closely related.

Theorem 1. *x is a rank efficient assignment $\Leftrightarrow \exists$ a valuation v such that x is supported by the v -rank-value mechanism.*

The intuition of this result comes from thinking about assignments in rank distribution space, that is mapping an assignment x to $(N^x(1), N^x(2), \dots, N^x(|\mathcal{O}|)) \in \mathbb{R}^{|\mathcal{O}|}$.²⁷ In rank distribution space, the proof of the theorem is similar to the proofs for the first and second welfare theorems. For instance, the backwards implication comes from rewriting the objective of the v -rank-value mechanism as²⁸

$$\sum_{k=1}^{|\mathcal{O}|-1} N^x(k) (v_k - v_{k+1}) + |\mathcal{A}| \cdot v_{|\mathcal{O}|}$$

²⁷Since the mechanism is individually rational, it will never assign any agents to their $(|\mathcal{O}| + 1)^{th}$ choice.

²⁸ $N^x(|\mathcal{O}|) = |\mathcal{A}|$ by definition, since all agents are assigned to some object (even if it is \emptyset).

By definition, if x is rank-dominated by \tilde{x} , then \tilde{x} will increase this objective. The forward implication comes from an argument that uses the non-standard polyhedral separating hyperplane theorem of McLennan (2002).²⁹ The details of the proof, we relegate to the Appendix.

The rank-value mechanism is not strategy-proof; in fact, in Section 7 we will show that no mechanism can be both rank efficient and strategy-proof. We will eventually address this shortcoming, but for now we merely state it as fact. Also note that linear programming mechanisms have been previously observed in the field (Roth 1991). In the next section, we will present other characterizations of rank efficiency and relate them to how Teach for America assigns teachers. Briefly, we mention that since rank-value mechanisms are ordinally efficient, their existence in the field helps to explain the puzzling absence of simultaneous eating mechanisms. Instead of ordinal efficiency most broadly, institutional evolution seems to have chosen the rank efficiency refinement.

5 Interpreting rank efficiency

Rank efficiency has several interpretations; going through them helps to make clear where it is present in the field. We will conclude the section by discussing the interpretations in the context of Teach for America’s system for assigning teachers.

5.1 Scoring interpretation

Rewriting the objective of the rank-value mechanisms as $\sum_k v_k \cdot [\sum_a \sum_o \mathbf{1}_{\{r_a(o)=k\}} \cdot x_{ao}]$ yields an easy interpretation. Score an assignment by giving it v_1 points for every agent (in expectation) who gets his first choice, v_2 points for every agent who gets his second choice, and so on. Then, look for the feasible assignment with the biggest score. This is an easy to understand explanation for non-economists, and it can have some justification depending on where the valuation v came from. For instance, say that we are assigning teachers to schools, but with some probability, a teacher will turn down the offer, even though he ranked it above \emptyset . Further, say that he is more likely to refuse an offer for a less preferred job than for more preferred job. If we view v_k as the probability that an agent will accept the offer given that the job was his k^{th} choice, then the rank-score mechanism is merely maximizing the expected number of agents who accept the offer. In fact, as we will discuss at the end of the section, this

²⁹The non-standard theorem is necessary to ensure that v_k is strictly larger than v_{k+1} for all k .

example very closely parallels a problem faced by Teach for America and is ostensibly one of their major motivations for using a rank-value mechanism.

5.2 Ex ante welfare interpretation

Let agent a 's von Neumann-Morgenstern utility of object o be $u_a(o)$. This is not a statement about “how many utils” o is worth to a , but rather a statement about how a values getting o with certainty relative to a lottery where, with probability p , she gets her most preferred object, and with probability $1 - p$, she gets her least preferred (possibly unacceptable) object. Using the affine degrees of freedom associated with von Neumann-Morgenstern utilities, set the utility of the top choice to 1 and the bottom (possibly unacceptable) choice to 0, for all agents.

When evaluating policy, it is often helpful to look at things from behind the veil-of-ignorance, that is, from the perspective of a fictional agent with no preferences of his own, who knows that he will randomly become one of the agents in the assignment market, inheriting her preferences. From this “original position”, Harsanyi (1975) proposes that a rational agent should act to maximize his expected utility. Assuming the von Neumann-Morgenstern axioms for our fictional agent, this gives an expected utility representation over agent-object pairs (Harsanyi 1955, 1986). Functionally, this representation is a weighted sum of the expected utilities of the agents, that is, it is a social welfare function, $\mathcal{W} = \sum_a \alpha_a \sum_o u_a(o)$.

The u_a 's encode how a values gambles over objects; this is an objective statement that can be falsified. The α 's, however, encode how our fictional agent will evaluate decisions such as whether he would rather get allocation x as agent a' with certainty or take a 50/50 gamble over getting allocation x as either agent a or agent a'' . Since our agent in the original position is only a thought experiment about justice, the α 's encode moral judgments that cannot be falsified.

Now that we have laid the groundwork for interpersonal comparison of utility, we can interpret the rank-value mechanism. The valuation v is an assumption that says that all agents will value gambles over objects in the same way (modulo differences in ordinal ranking). The fact that all agents are weighted the same in the objective of the rank-value mechanism means that we place the same social value on any agent getting a k^{th} choice. Of course, while reasonable, this is not the only assumption one could make. A more general version of the rank-value mechanism, which would allow for a policy-maker to make whatever assumption about interpersonal utility comparison he likes, is discussed in Section 8.1.

To be more precise, we could ask when a planner's assumptions make veil-of-ignorance welfare maximization equivalent to the rank-value mechanism. Let v_{ao} be agent a 's von Neumann-Morgenstern utility value for object o , and let F be a distribution over these values for all agents and objects. F encodes the planner's beliefs once he has seen the submitted ordinal preferences, as well as his assumption about how agents' utilities compare. Note then that it is without loss of generality to index v by agent, a , and rank, k . Given F , we assume that the policy-maker seeks to maximize $\int \sum_a \sum_o \sum_k v_{ak} \cdot \mathbf{1}_{\{r_a(o)=k\}} \cdot x_{ao} \cdot dF(v)$. The following proposition makes precise the assumptions the planner would need to make to justify running a rank-value mechanism.

Proposition 5. *Assume the planner wants to maximize welfare subject to assumptions about cardinal utilities and interpersonal utility comparisons that are encoded by the distribution F . If, relative to F , the unconditional expectation of v_{ak} is independent of a , then the planner can do so by running the \tilde{v} -rank-value mechanism, where $\tilde{v}_k = \int v_{ak} \cdot dF(v)$ (for any agent a).*

Proof. Changing the order of summation (and integration) in the welfare, we find

$$\begin{aligned} \sum_a \sum_o \sum_k \left[\int v_{ak} \cdot dF(v) \right] \cdot \mathbf{1}_{\{r_a(o)=k\}} \cdot x_{ao} \\ = \sum_a \sum_o \sum_k \tilde{v}_k \cdot \mathbf{1}_{\{r_a(o)=k\}} \cdot x_{ao} = \sum_a \sum_o \tilde{v}_{r_a(k)} \cdot x_{ao} \end{aligned}$$

□

Note that the condition of the Proposition can be interpreted either as a direct assumption, or as an informational limitation, i.e. the planner would have different beliefs for the agents if he knew more about them, but he doesn't.

Now, given the veil-of-ignorance interpretation we have just laid out, we can think of rank-efficiency as an ordinal adaptation of ex ante efficiency, that is, maximizing expected utility before agent's types are known.³⁰ The relationships between rank efficiency, ordinal efficiency, and ex post efficiency are expressed diagrammatically in Table 3.

Finally, we should mention what happens when other assumptions concerning ex ante welfare seem prudent. We could weight the different agents differently in the welfare function, which translates to a different assumption about how the social value

³⁰Of course, our model has no formal types; instead we can think of each agent as his own type.

of a k^{th} choice allocation compares across agents. This generalization is investigated further in Section 8.1. Another option would be to assume that we know the “type” of each agent and to make different assumptions about cardinal utility for different type agents. With such assumptions, we could certainly run a linear program that looks similar to a rank-score mechanism; however, we would not be guaranteed that it would yield a rank efficient assignment. Still, such a generalization is an interesting topic for future research.

5.3 Tough decisions interpretation

Motivating rank efficiency as ex ante efficiency, modulo a few assumptions, might be reassuring to theorists, but policy-makers tend to think in terms of deterministic assignments. If this is true, then the following claim might shed some light on why we don’t see the simultaneous eating mechanisms in the field.

Claim 9. On the domain of deterministic assignments, ordinal efficiency is equivalent to ex post efficiency.

Proof. On the deterministic domain, stochastic improvement is equivalent to ex post improvement. □

The same is not true, however, for rank efficiency

Claim 10. On the domain of deterministic assignments, rank efficiency implies ex post efficiency, but ex post efficiency does not necessarily imply rank efficiency.

Proof. We have already shown that rank efficiency \Rightarrow ordinal efficiency on a broader domain. To see the (\nRightarrow) part, consider \boxed{x} and $\bigcirc(x)$ from the leading example in Section 4. Both are ex post efficient deterministic assignments, but only one is rank efficient. □

	Ex post		Interim		Ex ante
Efficiency concept	Ex post efficiency	\Leftarrow \nRightarrow	Ordinal efficiency	\Leftarrow \nRightarrow	Rank efficiency
Mechanisms	Random serial dictatorships		Simultaneous eating mechanisms		Rank-score mechanisms

Table 3: Relationships between the ex post, ordinal and rank concepts of efficiency

The fundamental difference is that both ordinal efficiency and ex post efficiency are Paretian concepts, that is, they only look for improvements that make all agents weakly better off. Rank efficiency, on the other hand, is equipped with a method for looking at non-Pareto improvements and deciding whether the good outweighs the bad. In other words, rank efficiency can make “tough decisions” while ordinal and ex post efficiency cannot. Thinking about our efficiency concepts in terms of trade cycles can help us to make this intuition more precise. Formally,

Definition 7. A trade cycle on assignment \mathbf{x} is a sequence $\tau = ((a_1, o_1), \dots, (a_m, o_m))$ such that $x_{a_k o_k} > 0$ for each (a_k, o_k) in the sequence. **Implementing trade cycle τ on assignment \mathbf{x}** yields a new assignment, \tilde{x} such that, for all $k \in \{0, \dots, m\}$ (interpreting $k = 0$ and $k = m$ as the same index), $\tilde{x}_{a_k o_{k-1}} = x_{a_k o_{k-1}} + \min_k \{x_{a_k o_k}\}$, $\tilde{x}_{a_k o_k} = x_{a_k o_k} - \min_k \{x_{a_k o_k}\}$, and $\tilde{x}_{a_k o'} = x_{a_k o'}, \forall o' \in \mathcal{O} \setminus \{o_k, o_{k-1}\}$.

A trade cycle can be interpreted as a sequence of agents who hold probability shares of objects. Implementing a trade cycle means that a_1 gives some of his o_1 to a_2 and receives some o_m from a_m , and so on. Ordinal efficiency is characterized in terms of **improvement cycles**, that is, trade cycles in which $o_{k-1} \succ_{a_k} o_k$ for all k (strict for some k) (Bogomolnaia and Moulin 2001).

Proposition 6. x is ordinally efficient $\Leftrightarrow x$ is non-wasteful and admits no improvement cycles.

Rank efficiency has a similar characterization in terms of trade cycles in which some agents are made worse off but others are made better off. A valuation vector determines whether the good outweighs the bad. Formally,

Definition 8. A trade cycle τ on assignment x is a **v -rank-improving cycle** if
$$\sum_{\tau} (v_{r_{a_k}(o_{k-1})} - v_{r_{a_k}(o_k)}) > 0.$$

Theorem 2. x is rank efficient $\Leftrightarrow x$ is non-wasteful and \exists a valuation v such that x admits no v -rank-improving cycles.

The forward implication is obvious, but the reverse implication requires some subtlety. Intuitively, a change from assignment \tilde{x} to assignment \hat{x} can be decomposed into distinct trade cycles and acquisitions of unassigned objects. Now say that \hat{x} rank dominates \tilde{x} . Then, for any v , the objective of the v -rank-value mechanism must be bigger for \hat{x} than for \tilde{x} , which by linearity, means that at least one of these cycles or

acquisitions must improve the objective, which is exactly what the condition in the definition of a v -rank-improving cycle means.

Hence, the valuation vector v provides the rule that the v -rank-value mechanism uses to decide if trade cycles that aren't Pareto improving should be implemented. These "tough decisions" cycles are closely related to the concept of callousness brought up by Budish and Cantillon (forthcoming) and Budish (2009) in the context of multi-unit assignment. Consider the leading example of this section. Putting weight on the deterministic assignment \boxed{x} can be interpreted as the mechanism allowing 1 to callously take what he prefers without considering the greater good. Rank efficiency prevents such callousness, and in fact, we can interpret implementing v -rank-improving cycles as doing just this.

5.4 Teach for America: before and after

The story of Teach for America's recent redesign of their assignment system sheds light on how our characterizations of rank efficiency fit together. Teach for America (TFA) is a nationwide non-profit that puts selected college graduates into at-risk schools to teach. In 2011, TFA intends to assign around 8,000 teachers nationwide. The application process for potential TFA admits is arduous and extends over several rounds, including interviews and tests of teaching ability. Once the original pool of around 50,000 applicants has been reduced to about 20,000, the applicants submit a final round application, along with rankings over the regions to which TFA could potentially assign them.³¹ From these applications, the admissions team decides who should be made an offer *without* considering their regional preferences.³²

At this point, admissions are over and assignment begins. Note that admits get notification of admission and their assignment simultaneously. For this reason, TFA is very interested in making assignments that yield the most acceptances. As one might expect, the probability of an admit accepting TFA's offer is strongly dependent on whether that admit is getting a top choice region, just like the example in Section 5.1. Currently, TFA is considering modeling acceptance probabilities for use in their mechanism, which indicates that they take this interpretation of the rank-value mechanism seriously.

³¹There were 43 regions in the 2011-2012 assignment cycle.

³²When we suggested that assignment and admissions might be integrated, the idea was quickly rejected. Teach for America separates admission from assignment for two reasons. The first is that they feel they should admit the best candidates, and that it would be unfair to an unpopular region to lower the bar in finding its teachers. The second is that Teach for America feels that it can successfully persuade many admits into accepting their assignment, even if it isn't a top choice.

Before the 2010-2011 admissions cycle, TFA used an interesting system to assign the admits. First, a computer would choose an initial match. How exactly the computer did this is immaterial, but what is important is that the admissions team found the computer match to have an unacceptably bad rank distribution. To correct this, they would spend about a week in a conference room looking for trading cycles that could improve the match. Almost all of these were cycles where some admits were made worse off, while the rest were made better off. In other words, the cycles the assignment team implemented looked a lot like the “tough decisions” cycles described by the characterization of rank efficiency in Section 5.3.³³

We should mention that the reason we know the details of TFA’s assignment system is that TFA asked for a redesign of their assignment system during the 2010-2011 admissions cycle.³⁴ We observed the old process and tried to test our understanding by recreating it with an algorithm. Specifically, we tried to match their non-automated process with a rank-value linear program. This turned out to be quite successful. Of course, we worried about strategy-proofness and strongly cautioned TFA about the potential problems of manipulation, but upon seeing simulations that used random serial dictatorship, they decided that the costs of strategy-proofness were too high.³⁵ Although we had not intended the linear program to be what TFA should use, TFA pressed for it, not because it offered anything new, but because it accomplished in 30 seconds what had before taken two man-weeks. The perceived equivalence between the automated and non-automated processes used by TFA turns out to line up quite well with the theoretical results we have established. So, although TFA has only recently started using a rank-score mechanism to run their match, the non-automated system they used in the past also seems to have been a rank-efficient mechanism.

6 The potential gains of the rank-score mechanism

As mentioned previously, both the ordinally efficient and rank efficient mechanisms are generically non-strategy-proof. This could be a major problem, as agent manipulation of reported preferences could mean that, while assignments look good relative

³³TFA does not allow admits to rank \emptyset , and the number of admits was equal to the number of positions. In this environment, all feasible assignments are non-wasteful.

³⁴Through out this paper, I (Featherstone) have used the formal we, which has led to an ambiguity in this paragraph. The TFA redesign was jointly undertaken with Al Roth.

³⁵The other part of their reasoning was a strong belief that TFA admits wouldn’t or couldn’t game the system. They are geographically separated, only rank the regions once, and know little about the relative popularities of the regions. We formalize a version of this line of reasoning in Section 7.

to the submitted preferences, they could be quite bad relative to the true preferences. Backing away from strategy-proofness, then, is a risky move for a market designer. In this section, we consider a simple empirical exercise that sheds light on this risk. Say we have the true preferences from an assignment market. We can run uniform random serial dictatorship, the baseline strategy-proof mechanism, to get an idea of how efficient a match we can get without sacrificing strategy-proofness. Then, we can look at the counterfactual under which we run either probabilistic serial, or a rank-value mechanism. We can think of the efficiency of these counterfactuals as an upper bound on what we could gain. Of course, manipulation could completely undermine these gains, but if we don't see a significant difference under the truth-telling counterfactual, then stepping back from strategy-proofness is not even worth considering.

6.1 Assigning students to overseas programs at Harvard Business School

The data we will use to run our counterfactual exercise comes from the 2011 overseas program match at Harvard Business School (HBS). As part of its core curriculum, first year MBAs at HBS must participate in an overseas program.³⁶ They are assigned to a foreign company and remotely work on a project with that company during the first semester. The program culminates in a two-week trip over the winter break in which the MBA will present her work in person and be given the opportunity to make foreign business contacts.

In 2011, at the beginning of their first semester at HBS, 900 MBAs were asked to rank the 11 different countries to which they could be assigned.³⁷ The mechanism we³⁸ used to match the students was strategy-proof, so we feel comfortable in taking the preferences submitted to it as truthful. We should mention that preferences with indifferences were elicited from the MBAs. The strategy-proof adaptation of random serial dictatorship that we used is briefly described in the Appendix. In the main text, we will keep things simple by randomly breaking student indifferences and considering only strict preference assignment mechanisms. All results are qualitatively the same regardless of the method we use to break student indifferences, and regardless

³⁶At HBS, it is known as the FIELD 2 program.

³⁷Students only rank countries; once they are assigned to a country, the company they get is administratively assigned without their input.

³⁸Again, the fact that this paper is narrated using the formal we causes ambiguity. Al Roth and I redesigned the HBS match together.

of whether be break indifferences. Results from these alternate specifications are included in the Appendix.

6.2 The data

The analysis in this section is borrowed from Featherstone and Roth (2011). Figure 1 shows the rank distribution we get when we run uniform random serial dictatorship, probabilistic serial, or the v -rank-value mechanism, where the valuation vector is

$$v = (100, 80, 50, 35, 15, 10, 5, 3, 2, 1, 0.5)$$

The first thing to notice is that the rank distribution from probabilistic serial is essentially identical to that of uniform random serial dictatorship. The underlying assignments are also very similar, which means that there is little reason to favor the more complicated probabilistic serial mechanism over the simpler uniform random serial dictatorship. Che and Kojima (2010) show that, theoretically, the random assignment generated by the probabilistic serial mechanism asymptotically converges to the random assignment generated by uniform random serial dictatorship in the large market limit. Often with asymptotics, however, it is hard to know how large is large enough. In the case of HBS, 900 students and 11 countries seems to hit the mark.

Perhaps more striking is how well the v -rank-score mechanism performs. The number of students who get their first or second choice is increased by more than 15% when we move from the uniform random serial dictatorship to the rank-score mechanism. In absolute terms, this is about 120 students. So the gains from moving to a rank-value mechanism might indeed be large enough to justify backing away from strategy-proofness. Again though, the figures we are looking at are upper bounds. It could be that manipulation completely undermines these gains, or even makes the rank-value mechanism perform worse than random serial dictatorship.³⁹ We are merely making the case that backing away from strategy-proofness for the gains given by probabilistic serial does not make sense for the HBS match, but that doing so for the gains given by the rank-value mechanism might well be worth it.

³⁹Note that the work we do in Section 7 indicates that in at least some environments, we expect truth-telling.




Figure 1: Rank distributions for our mechanisms

6.3 Strategy-proofness versus efficiency

The previous exercise illustrates that the concept of rank efficiency and its associated rank-value mechanisms fit well with previous literature that has considered the costs of strategy-proofness, as first discussed by Erdil and Ergin (2008) and Abdulkadiroğlu, Pathak and Roth (2009) in the context of stable matching. Paralleling Azevedo and Leshno (2010), it remains an open question whether there exists an equilibrium of a rank-value mechanism that actually yields a worse rank distribution than uniform random serial dictatorship.

Briefly, we mention a few papers that consider strategy-proofness versus efficiency in the context of non-stable assignment. Mostly this has focused on how the Boston mechanism (Abdulkadiroğlu and Sönmez 2003*b*) can outperform random serial dictatorship. Featherstone and Niederle (2011) show this in the context of a truth-telling equilibrium, while Abdulkadiroğlu, Che and Yasuda (2011) focus on a non-truth-telling equilibrium. Both papers rely on stylized assumptions about how preferences are distributed. The present paper differs in two ways. First, it offers a new approach to get at cardinal utility in the context of an ordinal mechanism: make reasonable assumptions and maximize welfare accordingly. Second, in past literature, the leading contender with random serial dictatorship has been the class of Boston mechanisms. This paper offers a new mechanism that is based on the well-established empirical fact that policy-makers care about rank distributions.⁴⁰ As we have begun to show in this section, the rank-value mechanisms might be an important part of the discussion about the costs of strategy-proofness in ordinal assignment markets.

7 Incentives for truth-telling

One way that a rank-value mechanism could outperform random serial dictatorship in spite of its non-strategy-proofness is to have a manipulating equilibrium in which the efficiency gains don't disappear, much like in Abdulkadiroğlu, Che and Yasuda (2011). This sort of approach, while interesting, is not what we will pursue. Instead, we will look for environments in which truth-telling can be supported in equilibrium. First, however, we will show that rank efficiency and strategy-proofness are theoretically incompatible.

⁴⁰The Boston mechanism has been axiomatized. See Kojima and Ünver (2011).

7.1 An impossibility result

A mechanism is called **strategy-proof** if the allocation it gives to any agent when she truthfully reveals her ordinal preference stochastically dominates the allocation it would give her if she revealed anything else. A mechanism is called **weakly strategy-proof** if the allocation it gives to an agent when she deviates from truth-telling never strictly stochastically dominates what it gives her when she truthfully reveals. Another way to phrase the difference between these two concepts is in terms of the cardinal preferences that rationalize the true ordinal preferences. Strategy-proof means that truth-telling is a dominant strategy regardless of the rationalizing cardinal utilities. Weakly strategy-proof means that for any beliefs, there exist rationalizing cardinal utilities that make truth-telling a best-response.

Theorem 3. *No rank-efficient mechanism is strategy-proof. In fact, no rank efficient mechanism is even weakly strategy-proof.*

Proof. Consider a four agent, four object example.

$$\begin{aligned} 1 &: b \succ e \\ 2 &: b \succ c \succ e \\ 3 &: c \succ d \succ e \\ 4 &: d \succ b \end{aligned}$$

The unique⁴¹ rank-efficient allocation is $\{(1, e), (2, b), (3, c), (4, d)\}$. So under any rank-efficient mechanism, 1 must be assigned to e . Now, consider what happens if 1 deviates from truth-telling and submits $b \succ c \succ d \succ e$ instead. Now, the unique rank-efficient allocation is $\{(1, b), (2, e), (3, c), (4, d)\}$. Hence, under any rank-efficient mechanism, 1 gets b with the deviation, which first-order stochastically dominates the certain e he would get with the truth. \square

Intuitively, someone has to get stuck with e . Under the true preferences, we pay the smallest price for sticking 1 with e . When he submits $b \succ c \succ d \succ e$ instead, he is essentially exaggerating about how bad e is for him, forcing the mechanism to give it to 2 instead. Still, this example required quite a bit of specific knowledge. We might wonder whether truth-telling is a natural response in environments with less information.

⁴¹For an example of how to show that an assignment is the unique rank efficient assignment, see Claim 4 above.

7.2 A possibility result

Before presenting our result about when we might expect truthful preference revelation to a rank-value mechanism, we must first lay a bit of groundwork. Following Roth and Rothblum (1999), for some preference profile \succsim , define $\succsim^{o \leftrightarrow o'}$ to be the same preference profile, except all agents have switched the objects o and o' in their ordinal rankings. Similarly, let $x^{o \leftrightarrow o'}$ denote the (possibly infeasible) assignment where everyone assigned to o in x is reassigned to o' and vice-versa. Finally, define $(\tau, q)^{o \leftrightarrow o'}$ to simultaneously switch the capacities of o and o' , while redefining the tiebreaker, $\tau(x)^{o \leftrightarrow o'} \equiv \tau(x^{o \leftrightarrow o'})$. Note that $x^{o \leftrightarrow o'}$ will be feasible if the capacities are switched to $q^{o \leftrightarrow o'}$. An agent a 's beliefs are then summarized by a vector of random variables that range over potential ordinal preferences of the other agents (\succsim_{-a}), capacities of the objects (\mathbf{q}), and tie-breaker numberings, ($\boldsymbol{\tau}$). We treat \succsim_a , \mathcal{A} , and \mathcal{O} as known.

Definition 9. An agent's beliefs are **$\{o, o'\}$ -symmetric** if the distributions of $(\succsim_{-a}, \mathbf{q}, \boldsymbol{\tau})$ and $(\succsim_{-a}, \mathbf{q}, \boldsymbol{\tau})^{o \leftrightarrow o'}$ coincide. If the beliefs are $\{o, o'\}$ -symmetric for all $o, o' \in \mathcal{O} \setminus \{\emptyset\}$, then we simply call the beliefs **symmetric**.

One interpretation of this condition is that it represents a very symmetric environment. The interpretation we favor, however, is that it represents an environment in which agents have very little specific information, that is, an environment in which nothing really distinguishes one object from the other. Even when this is not globally true, it is often the case that preferences are tiered and that, within a tier, beliefs are close to symmetric. Note that this condition is well defined even for preferences that have indifferences. Under this informational assumption, we are guaranteed that switching the order of two objects in the submitted preference is not profitable. Formally,

Proposition 7. *Under a rank-value mechanism, if agent a 's beliefs are $\{o, o'\}$ -symmetric, and $o' \succ_a o$, then for a , the allocation under any submitted preference that declares $o \succ_a o'$ is weakly stochastically dominated by the allocation of a submitted preference that does not. If the beliefs are symmetric, then this holds for all $o, o' \in \mathcal{O} \setminus \{\emptyset\}$.*

The proof for the theorem leverages the symmetry assumption to show that for every state of the world in which the lie is profitable, there is at least one equally likely state of the world in which it is not. This result is stated in a way that allows for indifference, but from this point on, for expositional ease, we will limit ourselves to assignment markets where all preferences are strict. For the interested reader,

the more general results for indifferences are in the Appendix. Now, an immediate corollary of Theorem 7 is that if an agent has strict preferences and must rank all objects, then he does best to truth-tell if his beliefs are symmetric. Formally,

Theorem 4. *Under a rank-value mechanism, if an agent’s beliefs are symmetric, his preferences are strict, and he is required to rank all objects, then the allocation he receives under truth-telling weakly stochastically dominates the allocation he receives under any other strategy.*

Corollary to Theorem 4. *Under the conditions of Theorem 4, all agents truth-telling is an equilibrium.*

Although the requirement to rank all objects may seem unrealistic, it turns out that it is quite prevalent in the field. Consider public school assignment. Students are not required to rank all schools, but if they cannot be assigned to a school they ranked, they are generally given an administrative assignment.⁴² In this sort of situation, failing to rank all schools is equivalent to the student saying to the school district, “beyond what I ranked, you can fill out the rest of my rank-order list for me.” Teach for America also does not allow agents to rank the null object.

To understand manipulation when agents are not required to rank all objects, we must first lay a bit more groundwork, as we have reached the end of what we can do without resolving some of the ambiguity concerning the rank functions that we mentioned near the end of Section 2. First, define a rank scheme to be **upward-looking** if $r_a(o)$ is independent of the preferences among objects to which o is strictly preferred. Now, consider how we deal with the outcome of being unmatched in the case of an agent whose true preferences are $a \succ b \succ \emptyset$. One natural alternative is to treat \emptyset like any other object, that is $r(a) = 1$, $r(b) = 2$, and $r(\emptyset) = 3$. Often, however, policy-makers like to think of being unmatched as a distinct worst element in the overall rank distribution; that is, $r(\emptyset) = |\mathcal{O}|+1$, regardless of how many objects the agent declared acceptable. To help in thinking about this, consider inserting a new indifference class, I , which contains o , into its own new acceptable indifference class at the bottom of a preference \succsim_a (but right ahead of \emptyset). Call this new preference \succsim_a^{+I} . If $r_{\succsim_a^{+I}}(o) = r_{\succsim_a}(\emptyset)$, then we call the ranking scheme **unmatched-neutral**, as \emptyset is treated just as any other object. If $r_{\succsim_a^{+I}}(\emptyset) = r_{\succsim_a}(\emptyset)$, then we call the ranking

⁴²We know for certain that this is the policy in the San Francisco Unified School District, where we (I and Atila Abdulkadiroğlu, Muriel Niederle, Parag Pathak, and Al Roth) spent a year redesigning the assignment system; however, it is hard to imagine a school district that doesn’t work in this way, as all students legally must be able to claim *some* seat in a public school, even after the match has run.

scheme **unmatched-distinct**. In this language, the example where we treat \emptyset just like any other object is unmatched-neutral, and the example where we always give \emptyset the lowest rank possible is unmatched-distinct.

Now, we define two important classes of manipulations. Call a submitted preference \succsim'_a a **truncation** of the true preference \succsim_a if there is some object $o \succsim_a \emptyset$ such that $\succsim_a^{o \leftrightarrow \emptyset}$ is equivalent⁴³ to \succsim'_a . Similarly, call a submitted preference \succsim'_a an **extension** of the true preference \succsim_a if there is some object $\emptyset \succsim_a o$ such that $\succsim_a^{o \leftrightarrow \emptyset}$ is equivalent to \succsim'_a . Note that under these definitions, truth-telling is both a truncation and an extension.⁴⁴

Theorem 5. *If an agent's beliefs are symmetric, his preferences are strict, and the rank scheme is upward-looking and...*

- *...unmatched-distinct, then his allocation from any other strategy is weakly stochastically dominated by his allocation from playing some truncation.*
- *...unmatched-neutral, then his allocation from any other strategy is weakly stochastically dominated by his allocation from playing some extension.*

Note that we can't leverage both parts of the theorem simultaneously to find a condition for truth-telling.

Claim 11. A ranking scheme cannot be both unmatched-neutral and unmatched-distinct.

Proof. If it is unmatched-neutral, then $r_{\succsim_a + I}(p) = r_{\succsim_a}(\emptyset)$, which means that $r_{\succsim_a + I}(\emptyset) > r_{\succsim_a}(\emptyset)$, contradicting the assumption that the ranking scheme was unmatched-distinct. \square

Truncations are a familiar class of mechanisms,⁴⁵ and the intuition for why they can be profitable is simple. When the ranking scheme is unmatched-distinct, a truncation can be interpreted as an agent threatening the linear program: “give me something I like, or pay the price of giving me \emptyset ”. When the ranking scheme is unmatched-neutral, this threat no longer works, as truncating actually relaxes the optimization. So why extensions? Intuitively, consider an agent whose preference is $a \succ b \succ \emptyset$ in a

⁴³Since we are dealing with an individually rational mechanism, how things are listed below \emptyset does not matter.

⁴⁴In the definitions, we see this when $o = \emptyset$.

⁴⁵The paper in which our informational assumption was first used, Roth and Rothblum (1999), found that, in symmetric environments, agents on the proposed-to side of a deferred acceptance mechanism can always do just as well with a truncation.

market with 100 other agents whose preferences are $a \succ b \succ c \succ d \succ e \succ f \succ \emptyset$. If the agent doesn't extend, the linear program will assign him to \emptyset , since doing so is worth $v_3 - v_7$ more than the alternative of putting some other agent in \emptyset .

So even in low information environments, agents can gain, but it is notable that they gain by using a class of strategies very similar to those that can be used to game two-sided deferred acceptance in a low information environment. In the lab, Featherstone and Mayefsky (2011) are able to show that subjects fail to truncate under a deferred acceptance mechanism, leaving a significant amount of money on the table by doing so. It would be interesting to see if a similar result held for rank-value mechanisms,⁴⁶ and in fact, the experimental and computational work of Ünver (2001) and Ünver (2005) indicate that this is a real possibility.

8 Connection to competitive equilibrium mechanisms

A common idea for assignment is to harness the power of the market by creating a pseudomarket, endowing all agents with a budget of fiat⁴⁷ money, and calculating a competitive equilibrium. The first paper to operationalize this idea was Hylland and Zeckhauser (1979). In this section, we will show that rank efficiency has a strong connection to pseudomarket mechanisms. Before we do this, however, we need to expand our conceptual framework. First, we will introduce a generalization of rank efficiency in which agents are weighted differently in the rank distribution. Then, we will show that even when we consider the union of all assignments that are rank efficient relative to *some* weights, the resultant efficiency concept remains a refinement of ordinal efficiency.

8.1 Generalizing rank efficiency

Until now, we have weighted all agents equally in the rank distribution, which is equivalent to saying that the utility that any agent gets from a k^{th} choice is directly

⁴⁶The theorems we just proved can help us to design experiments that deal with the large strategy space of ordinal mechanisms, by allowing us to focus only on truncations or extensions. For a taste of how this might work, see Featherstone and Mayefsky (2011), which uses the “truncations are exhaustive” result of Roth and Rothblum (1999) to yield a more analytically tractable design when thinking about deferred acceptance mechanisms.

⁴⁷See Footnote 10 for a definition in this context.

comparable to the utility that any other agent gets from a k^{th} choice.⁴⁸ Now, we step back from this assumption by allowing for a general weighting scheme in the rank distribution. Formally, define the **rank distribution of assignment x with respect to weights α** (or the α -rank distribution for short) to be

$$N_{\alpha}^x(k) \equiv \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot \mathbf{1}_{\{r_a(o) \leq k\}} \cdot x_{ao}$$

$N_{\alpha}^x(k)$ is just the weighted expected number of agents who get their k^{th} choice or better under assignment x .

Definition 10. A feasible random assignment x is **α -rank dominated** by another feasible assignment \tilde{x} if the α -rank distribution of \tilde{x} first-order stochastically dominates that of x , that is, $N_{\alpha}^{\tilde{x}}(k) \geq N_{\alpha}^x(k)$ for all k (strict for some k). A random assignment is called **α -rank efficient** if it is not α -rank-dominated by any other feasible assignment.

The mechanisms that correspond to α -rank efficiency are a similarly modified version of the v -rank-value mechanisms.

Definition 11. The **rank-value mechanism with respect to weights α and valuation v** (or the **(α, v) -rank-value mechanism** for short) maps agent rank orderings to a maximizer of the following linear program.

$$\begin{aligned} \max_x \quad & \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao} \\ \text{s.t.} \quad & \sum_a x_{ao} \leq q_o, \quad \forall o \in \mathcal{O} \\ & \sum_o x_{ao} \leq 1, \quad \forall a \in \mathcal{A} \\ & x_{ao} \geq 0, \quad \forall o \in \mathcal{O} \\ & \quad \quad \quad \forall a \in \mathcal{A} \end{aligned}$$

We say that the assignments in the $\arg \max$ are **supported by the (α, v) -rank-value mechanism**.

Everything we have proved thus far is easily adapted to any given vector of weights α . We can think about the α vector as an assumption about how we should compare the utilities of different agents, as discussed in Section 5.2.

⁴⁸The interpersonal utility comparison is actually more subtle than this; see Section 5.2 for a discussion.

8.2 Supportable rank efficiency

Of course, we might not want to make any assumption about this. That is, we might be perfectly happy to make assumptions about agent’s cardinal utilities, but not about how those utilities compare. The following concept makes this idea precise.

Definition 12. An assignment x is called **supportably rank-efficient** if there exists an α such that x is α -rank-efficient.

Claim 12. Supportable rank efficiency implies ordinal efficiency.

Proof. Supportable rank efficiency implies, for some α , α -rank efficiency, which in turn, for any α , implies ordinal efficiency \square

At this point, the natural question whether supportable rank efficiency is equivalent to ordinal efficiency. It turns out that this is not the case.

Theorem 6. *Not all ordinally efficient allocations are supportably rank-efficient.*

The intuition here is rooted in appropriate generalization of the “tough decisions” trade cycles of Section 5.3. For an assignment to be supportably rank efficient, there must be one (α, v) pair such that all trade cycles fail to be (α, v) -rank improving.⁴⁹ This means that every trade cycle has some inequality associated with it. The proof of the theorem starts with an ordinally efficient assignment and then defines enough trade cycles to where the corresponding set of inequalities has no solution. There is nothing special about the counterexample we choose; as an assignment grows large, it has many trade cycles, so unless special provision is made to keep them all from being (α, v) -rank improving relative to some common (α, v) , it is no surprise that the system of inequalities cannot be satisfied. Although we have not proven it, the proof suggests a natural conjecture: as the market grows, the size of the set of supportably rank-efficient assignments becomes smaller when compared to the size of the set of ordinally efficient assignments. Also, note that there is no reason that probabilistic serial assignments are special; in fact, the proof to Theorem 6 uses a probabilistic serial assignment to get the ordinally efficient assignment that forms the counterexample.

Corollary to the proof. *Probabilistic serial does not always yield supportably rank-efficient assignments.*

⁴⁹Formally, in Definition 8, replace $\sum_k (v_{r_{a_k}(o_{k-1})} - v_{r_{a_k}(o_k)}) > 0$ with $\sum_k \alpha_{a_k} \cdot (v_{r_{a_k}(o_{k-1})} - v_{r_{a_k}(o_k)}) > 0$.

This result might seem especially surprising in the light of McLennan (2002) and Manea (2008), who both find (using non-constructive and constructive methods, respectively) that any ordinally efficient assignment is ex ante efficient relative to some set of cardinal preferences that rationalize the ordinal preferences. The difference here is that, as we showed in Section 5.2, the concept of rank-efficiency imposes the added constraint that all agents must share the same cardinal utility across ranks, that is, for any ranks j and k , the ratio of agent a 's utility for her j^{th} ranked object and her k^{th} ranked object must be the same for all agents a .

8.3 Competitive equilibrium mechanisms

A commonly considered idea in random assignment is to take submitted preferences, give all agents a budget of fiat⁵⁰ money and calculate a competitive equilibrium. This class of pseudomarket mechanisms for cardinal preferences was first considered in Hylland and Zeckhauser (1979). Since we are dealing with ordinal preferences, we assume a valuation vector and adapt Hylland and Zeckhauser (1979) to our ordinal setting.

Definition 13. The prices and random assignment (\tilde{x}, \hat{p}) , form a **budget equilibrium with respect to valuation v and budgets B** if

$$\begin{aligned} \tilde{x}_a &\in \arg \max_{x_a \geq 0} \sum_o v_{r_a(o)} \cdot x_{ao} \\ \text{s.t.} \quad &\sum_o \hat{p}_o \cdot x_{ao} \leq B_a, \forall a \\ &\sum_o x_{ao} \leq 1 \end{aligned}$$

and

$$\begin{aligned} \tilde{x}_{ao} &\in \arg \max_{x_{ao} \geq 0} \sum_a \sum_o \hat{p}_o \cdot x_{ao} \\ \text{s.t.} \quad &\sum_a x_{ao} \leq q_o \end{aligned}$$

A **(B, v) -budget mechanism** maps ordinal preferences to an assignment that is part of a (B, v) -budget equilibrium.

⁵⁰See Footnote 10 for a definition in this context.

We can think of the first optimization as the agent's problem, and the second as the producer's problem.

Remark 5. The producer's problem is equivalent to $\sum_a x_{ao} \leq q_o$ and $\hat{p}_o > 0 \Rightarrow \sum_a x_{ao} \leq q_o$, for all $o \in \mathcal{O}$.

The resemblance to the familiar Walrasian equilibrium is enough to make us think that some results resembling the welfare theorems must be present.⁵¹

Claim 13. A (B, v) -budget mechanism's assignment must be ordinally efficient.

Proof. Say that it isn't. Then there is a Pareto improvement, which cannot be, since the pseudomarket yields an ex ante Pareto efficient allocation relative to the cardinal utilities v , by a slight adjustment of the first welfare theorem argument presented in Hylland and Zeckhauser (1979). \square

The previous claim was a version of the first welfare theorem; however, we won't find a version of the second. To get both welfare theorems, it turns out that supportable rank efficiency is the correct efficiency concept. Put less technically, the budget mechanisms are the supportably rank efficient mechanisms.

Theorem 7. x is supportably rank efficient $\Leftrightarrow \exists(B, v)$ such that x is supported by the (B, v) -budget mechanism.

⁵¹Note that the (α, v) -rank-value mechanism can be decentralized to look like a Walrasian equilibrium as well.

Definition 14. Prices and an assignment, (x^*, \hat{p}) , form a **discount equilibrium with respect to weights α and valuation v** (or just (α, v) -discount equilibrium) if

$$x_a^* \in \arg \max_{x_a \geq 0} \sum_o \left(v_{r_a(o)} - \frac{\hat{p}_o}{\alpha_a} \right) \cdot x_{ao}, \forall a$$

$$\text{s.t.} \quad \sum_o x_o \leq 1$$

and

$$x_{ao}^* \in \arg \max_{x_{ao} \geq 0} \sum_a \sum_o \hat{p}_o \cdot x_{ao}$$

$$\text{s.t.} \quad \sum_o x_{ao} \leq q_o$$

Proposition 8. x is α -rank efficient $\Leftrightarrow \exists \hat{p}, v$ such that (x^*, \hat{p}) form an (α, v) -discount equilibrium

This is a strange decentralization, as it distorts prices in a way that one might expect it to not yield any sort of efficiency. The proof of the Proposition, as well as some discussion on why a price-distorted equilibrium concept should yield any sort of efficiency can be found in the Appendix.

The intuition for this proof comes from considering the linear programming duals of the optimizations involved in the (α, v) -rank-value mechanism and the (B, v) -budget mechanism. If we start with a supportably rank efficient assignment x , then it must be that x is in the arg max of the (α, v) -rank-value mechanism for some (α, v) . We can keep the same valuation v in our budget equilibrium. The prices \hat{p} that will decentralize x are the shadow prices associated with the object quota constraints in the (α, v) -rank-value mechanism. To make sure that agent a can afford his allocation, x_a , we give him a budget $B_a = \sum_o \hat{p}_o \cdot x_{ao}$.

To go in the other direction, we start with an assignment x that is a (B, v) -budget equilibrium. Again, we can stick with the same valuation v , but we have to choose weights α such that x is in the arg max of the (α, v) -rank-value mechanism. To do this, we set α_a to the inverse of the shadow price of budget in agent a 's optimization.⁵²

An important corollary to the proof of this proposition is that, for a given valuation v and set of preferences \succsim , we have a natural mapping between weights α and budgets B . What's more if, with a mind toward procedural fairness, we run a v -rank-value mechanism, then the assignments in our arg max will be supported by a budget equilibrium in which budgets are not guaranteed to be equal. Similarly, if we start with equal budgets, then our assignment will be supported by an (α, v) -rank-value mechanism in which the weights are not guaranteed to be equal. In this sense, procedural fairness is fundamentally different in our two mechanisms.

9 Justice and rank-efficiency

Looking at the leading example from Section 4, one might object that agent 1 has been unfairly singled out. Envy-freeness is the idea that makes this intuition precise. Let an assignment be **strongly envy-free** if x_a weakly first-order stochastically dominates $x_{a'}$ relative to \succsim_a , for all $a, a' \in \mathcal{A}$. Let an assignment be **weakly envy-free** if there is no $a' \in \mathcal{A}$ such that $x_{a'}$ strongly first-order stochastically dominates x_a relative to \succsim_a , for all $a \in \mathcal{A}$. Finally, let an assignment be **envy-free** relative to some cardinal preferences if each agent weakly prefers his own allocation to all other agents' allocations.

Theorem 8. *No mechanism is rank efficient and even weakly envy-free*

Proof. Consider the abbreviated example from Section 4. (x) is the unique rank

⁵²We have to be a little careful when the shadow price is zero. This case is worked out in more detail in the Appendix.

efficient assignment, so any rank efficient mechanism must choose it. But 1 would prefer 2's allocation, regardless of the cardinal preferences that rationalize his ordinal preference. \square

Before we go any further, it behooves us to press a bit more into what this theorem means. Justice is a philosophically tricky subject, but a good start is to consider Dworkin's concept of equality of resources. The idea is rooted in the thought that no agent in the market should envy another agent's allocation. In addition to this, however, Dworkin adds a requirement of the sort of efficiency the market provides. After all, it is easy to imagine envy-free assignments that lack efficiency.⁵³ In Dworkin (1981), he literally suggests a cardinal version of this paper's budget equilibrium from equal budgets. In contrast to equality of resources, we have Harsanyian justice (Harsanyi 1975). Here, a just policy is chosen by an expected utility maximizer in the original position.⁵⁴ Under this conception, we might decide to have an envious agent if doing so were to help many other agents.⁵⁵ The "tough decisions" of Section 5.3 are part and parcel to Harsanyian justice.

Now that we have introduced these two concepts of justice, Theorem 8 comes into sharper focus. It says that in assignment markets, the fairness concepts of Dworkin and Harsanyi are mutually exclusive. This is no surprise in view of the proof of Theorem 7. If an assignment is implemented by a budget equilibrium from equal budgets, then it is also implemented by a rank-score mechanism with unequal weights. If an assignment is implemented by a rank-score mechanism with equal weights, then it is also implemented by a budget equilibrium with different budgets. So to move forward from Theorem 8, we must choose either Dworkin or Harsanyi. As discussed in Section 5.2, the rank-value mechanism with equal weights implements Harsanyian justice modulo a few assumptions. But what if we want Dworkin justice? Due to our ordinal setting, we won't be able to get strong envy-freeness, but budget equilibrium from equal budgets will get us weak envy-freeness (and envy-freeness, modulo a few assumptions).

Proposition 9. *If $\exists \hat{p}, v, B$ such that $B_a = B_{a'}$ for all $a, a' \in \mathcal{A}$ and (x, \hat{p}) is a (B, v) -budget equilibrium, then x is weakly envy-free and α -rank efficient for some α .*

⁵³For instance, assign all agents to \emptyset .

⁵⁴That is, from the perspective of a fictional agent who knows that they will become one of the agents in the market, uniformly and at random, and will inherit that agent's preferences and allocation.

⁵⁵Although Harsanyi's conception of justice is based on Rawls' veil-of-ignorance, the Difference Principle (i.e. insistence on a maximin objective from the original position) places Rawls' conception more in line with Dworkin's (Rawls 1972).

Further, x is envy-free relative to the cardinal preferences encoded by v .

Proof. Budget equilibria are supportably rank efficient, so x must be α -rank efficient for some α . Now, assume that $x_{a'}$ strongly first-order stochastically dominates x_a relative to \succsim_a . Then, for any v , agent a 's objective in the definition of the (B, v) -budget equilibrium must be higher at $x_{a'}$ than at x_a . But a could have afforded $x_{a'}$ by the equal budgets assumption. This contradicts optimality. Envy-freeness relative to v follows by similar logic. \square

Although this provides a partial workaround to the impasse we just discussed, weak envy-freeness is not a strong condition. The proposition only guarantees that there *exist* rationalizing cardinal utilities that make the assignment envy-free. Still, the exercise is very similar to how we justified the v -rank-value mechanism as maximizing social welfare, that is, in the absence of cardinal information, we make assumptions and push forward. In that sense, given that we are willing to make an assumption about rationalizing cardinal utilities, budget equilibrium from equal budgets is to envy-freeness what the rank-value mechanism is to ex ante efficiency. Finally, we should note that the probabilistic serial mechanism of Bogomolnaia and Moulin (2001) is strongly envy free, but as we showed in Section 8.2, it may not be ex ante efficient in a world where our beliefs about agents' cardinal preferences meet the condition of Proposition 5.⁵⁶

10 Conclusion

Rank efficiency is a natural concept that is used in the field, and under truth-telling, rank efficient mechanisms can yield significant efficiency gains. Unfortunately, rank efficiency is also theoretically incompatible with strategy-proofness. This situation is not unique: there is a growing body of literature that points to the fact that the costs of strategy-proofness can be quite high. When then should the market designer pay this cost? In the paper, we showed that truth-telling is an equilibrium in a stylized, low-information environment. Such a theorem gives us the important intuition that rank-value mechanisms are more likely to work when agents know little about the popularity of the objects. Unfortunately, we do not have results about equilibrium in other environments, and in fact, it seems likely that there exist environments in

⁵⁶Briefly, we note that the wedge between fairness and efficiency exists in the cardinal setting as well. Competitive equilibrium from equal incomes (Varian 1974, Hylland and Zeckhauser 1979) gives Pareto efficiency, but Negishi's theorem (Negishi 1960) does not guarantee that the corresponding weights in the planner's problem will be equal.

which the performance of rank efficient mechanisms will be unacceptably poor. Unfortunately, any such characterization is likely to again be limited to a stylized model, and characterizations that truly generalize across environments seem intractable. In addition, it is certainly plausible that in more complicated situations, equilibrium predictions will not be born out in the field. Thus, to understand when the costs of strategy-proofness are too dear, we will need to complement theoretical results with more empirical approaches, such as experiments and learning models. We are currently pursuing both of these avenues in parallel work: the computational agenda in Featherstone and Roth (2011), and the experimental agenda in an experiment similar in design to Featherstone and Mayefsky (2011).

Regardless of such evidence, it may be that, in the long run, in any environment in which a mechanism fails to admit a truth-telling equilibrium, agents will eventually converge to an equilibrium. Even if this is true, we argue that it is still a worthwhile exercise to consider how much the short run can be extended. Consider antibiotics. In the long run, it is inevitable that resistant pathogens will evolve; however, in the short-run, lives are saved. Less rhetorically, consider the HBS match discussed in Section 6. In the first year of the match, had we run a rank-efficient mechanism, it is quite plausible that agents would have truthfully revealed their preferences. We didn't do this, however, because we were worried about the MBAs learning to manipulate over time. Taking extra efficiency in the first year, at the expense of subsequent years, seemed like too much of a "smash-and-grab" strategy. But what if the learning process were slower, such that we could reap the efficiency gains for the first five years? In such a setting, the "smash-and-grab" strategy starts to sound like a good idea; in fact, we could even imagine running a rank efficient mechanism for four years and then switching to a strategy-proof mechanism as learning began to cause problems. Coming back to the antibiotics analogy, pharmaceutical research extends the short-run by manufacturing different types of antibiotics, so that the evolutionary process (so far) is unable to keep up. Thinking about ways to slow the speed with which agents learn to manipulate a non-strategy-proof mechanism seems like an interesting, and thus far, unpursued line of research. When the efficiency cost of strategy-proofness is high, as we have shown can be the case, this line of inquiry could prove quite fruitful.

In conclusion, the fact that we see linear programming mechanisms in the field could imply that rank-efficient mechanisms are successful in some situations, but it could also mean that policy-makers have implemented a bad policy. Market designers should either be correcting the impulse of policy-makers to depart from strategy-

proofness, or they should be sometimes suggesting that strategy-proofness is a small cost to pay for big efficiency gains. Which advice is correct will very likely be strongly dependent on the environment in which the mechanism is meant to serve, as well as the dynamics of the social learning process. Understanding when (if ever) market designers can feel comfortable suggesting a rank efficient mechanism thus remains an important open question.

References

- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 1998. “Random serial dictatorship and the core from random endowments in house allocation problems.” *Econometrica*, 689–701.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003*a*. “Ordinal efficiency and dominated sets of assignments.” *Journal of Economic Theory*, 112(1): 157–172.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003*b*. “School choice: A mechanism design approach.” *The American Economic Review*, 93(3): 729–747.
- Abdulkadiroğlu, Atila, Parag Pathak, and Alvin E. Roth.** 2009. “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the New York City High School Match.”
- Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda.** 2011. “Resolving conflicting preferences in school choice: The Boston mechanism reconsidered.” *The American Economic Review*, 101(1): 399–410.
- Azevedo, Eduardo M., and Jacob Leshno.** 2010. “Can We Make School Choice More Efficient?” *unpublished mimeo*.
- Berge, Claude.** 1957. “Two theorems in graph theory.” *Proceedings of the National Academy of Sciences of the United States of America*, 43(9): 842.
- Birkhoff, Garrett.** 1946. “Tres observaciones sobre el algebra lineal.” *Univ. Nac. Tucumán Rev. Ser. A*, 5: 147–151.
- Bogomolnaia, Anna, and Hervé Moulin.** 2001. “A new solution to the random assignment problem.” *Journal of Economic Theory*, 100(2): 295–328.
- Bogomolnaia, Anna, Rajat Deb, and Lars Ehlers.** 2005. “Strategy-proof assignment on the full preference domain.” *Journal of Economic Theory*, 123(2): 161–186.
- Budish, Eric.** 2009. “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes.” *University of Chicago, Booth School of Business, Unpublished mimeo*.
- Budish, Eric, and Estelle Cantillon.** forthcoming. “The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard.” *The American Economic Review*.
- Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom.** 2011. “Implementing Random Assignments: A Generalization of the Birkhoff-von Neumann Theorem.” *University of Chicago, Columbia University, Stanford University, unpublished mimeo*.
- Che, Yeon-Koo, and Fuhito Kojima.** 2010. “Asymptotic equivalence of probabilistic serial and random priority mechanisms.” *Econometrica*, 78(5): 1625–1672.

- Cook, William J., William H. Cunningham, William R. Pulleybank, and Alexander Schrijver.** 1998. *Combinatorial Optimization*. New York, NY: Wiley-Interscience.
- Dworkin, Ronald.** 1981. “What is equality? Part 2: Equality of resources.” *Philosophy and Public Affairs*, 10(4): 283–345.
- Egerváry, Jéno.** 1931. “Matrixok kombinatorius tulajdonságairól.” *Matematikai és Fizikai Lapok*, 38: 16–28.
- Erdil, Aytek, and Haluk Ergin.** 2008. “What’s the matter with tie-breaking? Improving efficiency in school choice.” *The American Economic Review*, 98(3): 669–689.
- Featherstone, Clayton R., and Alvin E. Roth.** 2011. “Could a Rank Efficient Mechanism Work? Empirical Evidence from a Harvard Business School Match.” *mimeo*.
- Featherstone, Clayton R., and Eric S. Mayefsky.** 2011. “Why Do Some Clearinghouses Yield Stable Outcomes? Experimental Evidence on Out-of-Equilibrium Truth-telling.” *mimeo*.
- Featherstone, Clayton R., and Muriel Niederle.** 2011. “School Choice Mechanisms under Incomplete Information: An Experimental Investigation.” *Harvard Business School, Unpublished manuscript*.
- Gale, David, and Lloyd S. Shapley.** 1962. “College admissions and the stability of marriage.” *The American Mathematical Monthly*, 69(1): 9–15.
- Hall, Philip.** 1935. “On representation of subsets.” *J. London Math. Soc.*, 10(2): 6–30.
- Harsanyi, John C.** 1955. “Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility.” *The Journal of Political Economy*, 63(4): 309–321.
- Harsanyi, John C.** 1975. “Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls’s Theory.”
- Harsanyi, John C.** 1986. *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge Univ Pr.
- Hopcroft, John E., and Richard M. Karp.** 1973. “An algorithm for maximum matching in bipartite graphs.” *Society for Industrial and Applied Mathematics J. Comput.*, 2: 225–231.
- Hylland, Aanund, and Richard Zeckhauser.** 1979. “The efficient allocation of individuals to positions.” *The Journal of Political Economy*, 293–314.

- Katta, Akshay-Kumar, and Jay Sethuraman.** 2006. “A solution to the random assignment problem on the full preference domain.” *Journal of Economic theory*, 131(1): 231–250.
- Kojima, Fuhito, and Mihai Manea.** 2010. “Incentives in the probabilistic serial mechanism.” *Journal of Economic Theory*, 145(1): 106–123.
- Kojima, Fuhito, and M. Utku Ünver.** 2011. “The “Boston” School-Choice Mechanism.” *unpublished mimeo*.
- Kuhn, Harold W.** 1955. “The Hungarian method for the assignment problem.” *Naval research logistics quarterly*, 2(1-2): 83–97.
- Manea, Mihai.** 2008. “A constructive proof of the ordinal efficiency welfare theorem.” *Journal of Economic Theory*, 141(1): 276–281.
- Manea, Mihai.** 2009. “Asymptotic ordinal inefficiency of random serial dictatorship.” *Theoretical Economics*, 4(2): 165–197.
- McLennan, Andrew.** 2002. “Ordinal Efficiency and the Polyhedral Separating Hyperplane Theorem.” *Journal of Economic Theory*, 105(2): 435–449.
- Munkres, James.** 1957. “Algorithms for the assignment and transportation problems.” *Journal of the Society for Industrial and Applied Mathematics*, 32–38.
- Negishi, Takashi.** 1960. “Welfare economics and existence of an equilibrium for a competitive economy.” *Metroeconomica*, 12(2-3): 92–97.
- NYC Department of Education.** 2009. “High School Admissions Press Release.”
- Papadimitriou, C.H., and K. Steiglitz.** 1998. *Combinatorial optimization: algorithms and complexity*. Dover Pubns.
- Rawls, John.** 1972. *A Theory of Justice*. Oxford, Clarendon Press.
- Roth, Alvin E.** 1984. “The evolution of the labor market for medical interns and residents: a case study in game theory.” *The Journal of Political Economy*, 991–1016.
- Roth, Alvin E.** 1991. “A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom.” *The American Economic Review*, 415–440.
- Roth, Alvin E., and Elliott Peranson.** 1999. “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design.” *The American Economic Review*, 89(4): 748–780.
- Roth, Alvin E., and Uriel G. Rothblum.** 1999. “Truncation strategies in matching markets – in search of advice for participants.” *Econometrica*, 67(1): 21–43.

- San Francisco Unified School District.** 2011. “Student Assignment Report.”
- Schrijver, Alexander.** 1998. *Theory of Linear and Integer Programming*. John Wiley & Sons Inc.
- Schrijver, Alexander.** 2003. *Combinatorial optimization*. Springer.
- Svensson, Lars-Gunnar.** 1994. “Queue allocation of indivisible goods.” *Social Choice and Welfare*, 11(4): 323–330.
- Svensson, Lars-Gunnar.** 1999. “Strategy-proof allocation of indivisible goods.” *Social Choice and Welfare*, 16(4): 557–567.
- Ünver, M. Utku.** 2001. “Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets.” *Journal of Economic dynamics and control*, 25(6-7): 1039–1080.
- Ünver, M. Utku.** 2005. “On the survival of some unstable two-sided matching mechanisms.” *International Journal of Game Theory*, 33(2): 239–254.
- Varian, Hal.** 1974. “Equity, envy, and efficiency.” *Journal of Economic Theory*, 9(1): 63–91.
- Zhou, Lin.** 1990. “On a conjecture by Gale about one-sided matching problems.” *Journal of Economic Theory*, 52(1): 123–135.

Appendix

A Ex post dominance expressed as state-by-state dominance

Consider a discrete probability space, $(\Omega, 2^\Omega, p(\cdot))$. Formally, a lottery over deterministic assignments can be thought of as a random variable X that maps each element of the state space, $\omega \in \Omega$, into a deterministic assignment. State-of-the-world ω occurs with probability $p(\omega)$, so the probability that a given deterministic assignment x is drawn is $p(x) = \sum_{\omega} \mathbf{1}_{\{X(\omega)=x\}} \cdot p(\omega)$. Hence, X represents a lottery over assignments in which a given deterministic assignment x is chosen with probability $p(x)$. If two lotteries can be embedded in some discrete probability space as random variables X and X' in such a way that all agents weakly prefer $X(\omega)$ to $X'(\omega)$ for all $\omega \in \Omega$, with strict preference for at least one agent in one state $\omega \in \Omega$, then we say that the lottery embedded as X **state-by-state ex post dominates** the lottery embedded as X' .

Claim 14. x is ex post efficient $\Leftrightarrow x$ has a lottery representation that cannot be state-by-state ex post dominated by any other lottery over feasible deterministic assignments

Proof. (\Leftarrow) Assume that the support of any lottery representation of ex post efficient assignment x contains at least one ex post dominated deterministic assignment. Replacing this element of the support with its ex post dominator yields the state-by-state dominator for every lottery representation of x , a contradiction.

(\Rightarrow) Assume that any lottery representation of x has a state-by-state dominator. Then one of these ex post efficient assignments must be ex post dominated, meaning that there is no representation of x that is a lottery over ex post efficient assignments, a contradiction. \square

B Random serial dictatorship with indifferences

Several papers have looked at serial dictatorships adapted for indifferences (Svensson 1994, 1999, Bogomolnaia, Deb and Ehlers 2005), but they have mainly focused on theoretical characterizations. In this section of the appendix, we will briefly describe the theoretical mechanism and then two methods by which it can be implemented.

Let π be a permutation of $(1, \dots, |\mathcal{A}|)$, and let $\rho_{\pi(k)}$ be the rank of $a_{\pi(k)}$'s most preferred objects that she can be assigned while still assigning agent $a_{\pi(k')}$ something that he ranked $\rho_{\pi(k')}$, for all $k' < k$. Define **serial dictatorship with indifferences relative to π** as the mechanism that yields any deterministic assignment that fulfills the rank guarantees we just recursively defined.

Basically, we can think of this as a random serial dictatorship in which agents are serially assigned an indifference class guarantee instead of an actual object. Define **random serial dictatorship with indifferences** analogously to the random serial dictatorship without indifferences defined in Section 3. Svensson (1994) proves that, even with indifferences, a deterministic assignment is ex post efficient if and only if it can be generated by serial dictatorship with indifferences relative to some ordering π . The extension to random assignments, an analog of Proposition 1 in the main text, is straightforward.

Proposition 10. *An assignment x is ex post efficient $\Leftrightarrow \exists$ a distribution over dictatorship orderings such that the corresponding random serial dictatorship with indifferences can yield x*

Also true, and perhaps obvious, is that the extension of random serial dictatorship to indifferences is strategy-proof.

Proposition 11 (Svensson 1994). *Random serial dictatorship with indifferences is strategy-proof, so long as the distribution over dictatorship ordering is fixed before agents reveal their preferences.*

A natural question at this point is “How does one run this mechanism practically?” Svensson (1994) offers an algorithm for calculating the opportunity sets of the agents as the dictatorship progresses, but not for calculating an actual assignment. We offer two more practical algorithms.

B.1 Integer programming approach

We can easily calculate an agent's rank guarantee with a simple linear integer program

$$\begin{aligned}
 \rho_{\pi(k)} = \max_x \quad & r_{a_{\pi(k)}}(o) \\
 \text{s.t.} \quad & \sum_a x_{ao} \leq q_o \\
 & \sum_o x_{ao} = 1 \\
 & x_{ao} \in \{0, 1\} \\
 & r_{a_{\pi(k')}}(o) \cdot x_{a_{\pi(k')}o} = \rho_{a_{\pi(k')}} \cdot x_{a_{\pi(k')}o}
 \end{aligned}$$

Any assignment in the arg max of the calculation of $\rho_{|\mathcal{A}|}$ is generated by serial dictatorship with indifferences relative to π . The advantage to this approach is that random serial dictatorship with indifferences can be easily calculated with off-the-shelf linear optimization software, such as CPLEX or the open-source GNU Linear Programming Kit (GLPK).

B.2 Augmenting paths approach

The downsides of the integer programming approach are that it uses black box software and that for large problems, the integer program might solve quite slowly. For these reasons, we introduce a more primitive algorithm that can be coded from scratch if necessary. Consider a bipartite graph connecting agents to objects (q_o different nodes for each object o) with two kinds of links – potential and realized. **Realized** links represent tentative allocations, while **potential** links represent allocations that we could implement if we so chose. At any stage in the algorithm we are about to present, the realized links will be a near maximum matching in which only one agent has no realized link, but does have potential links. To see if we can match our unmatched agent, we will need to look for an **augmenting path**, which is a path $(a_1, o_2, a_3, o_4, \dots, a_{n-1}, o_n)$ where (a_j, o_{j+1}) is a potential link, and (o_j, a_{j+1}) is a realized link. **Implementing an augmenting path** means switching all of its potential links to realized links, and vice-versa. This necessarily increases the number of agents in the realized link matching by one. In fact, Berge’s theorem (Berge 1957) tells us that a matching is a maximum matching if and only if it has no augmenting path. With this background, the algorithm can be laid out.

1. Initialize.
 - (a) Create the graph G_0
 - i. Put a potential link between all agents and objects.
 - ii. Pick any feasible deterministic assignment and put a realized link between every agent and objects that are assigned to each other.
 - (b) Set $i = 1$ and $k = 1$.
2. Let O_k be the set of objects in the k^{th} indifference class of agent $\pi(i)$. Create G_i^{temp} by deleting all links (implemented and potential) in G_{i-1} that connect to agent $\pi(i)$ and replacing them with potential links between agent $\pi(i)$ and every object in O_k .

3. If there is an augmenting path in G_i^{temp} , implement it, set $G_i = G_i^{temp}$, increment i , set $k = 1$, and go to Step 4. If not, then increment k and go to Step 2.
4. If $i \leq |\mathcal{A}|$, then go to Step 2. Otherwise terminate the algorithm.

When the algorithm terminates, the realized links in $G_{|\mathcal{A}|}$ represent a match is generated by serial dictatorship with indifferences relative to π .

The algorithm is simple. As we go through the agents, serially, we try to match them to the best indifference class we can, while preserving a feasible match for the rest of the agents. Berge’s theorem assures us that if we can’t find an augmenting path, then our agent in question cannot be matched to anything in the indifference class we are trying for. The only computationally difficult parts of this algorithm are finding the initial matching and searching for an augmenting path in a bipartite graph. Finding the initial matching can be accomplished via the Hopcroft-Karp algorithm (Hopcroft and Karp 1973), and finding an augmenting path is typically accomplished with a breadth first search. Algorithms for these subroutines are easily found in most textbooks on combinatorial optimization (Schrijver 2003, Cook et al. 1998, Papadimitriou and Steiglitz 1998).

C The HBS overseas match: indifferences and how they were broken

Again, most the results from this section are borrowed from Featherstone and Roth (2011). The HBS match, as mentioned in the main text, allowed for indifferences to be submitted. It was run using the strategy-proof extension of random serial dictatorship described in Section B. In the main text, however, we showed results for strict preferences, which we constructed by breaking the ties in each agent’s submitted preference. There are two reasonable ways to do this. One is to break the ties by independently replacing each indifference class with a uniform random draw from the sets of strict orderings over the objects in the class. We call this the uncorrelated tie-breaker. The other way we looked at was to break all indifferences in the alphabetical order of the regions, which was also the order in which the regions were listed on the website used to elicit preferences from the agents. We call this the correlated tie-breaker. The results from the correlated tie-breaker are presented in Figure 1 in the main text, while the results from the uncorrelated tie-breaker are presented in Figure 2 in the Appendix.




Figure 2: Results from the uncorrelated tie-breaker

We also present the actual results from the HBS match when we deal with indifferences instead of breaking them in Figure 3. Also included in the figure is the rank distribution from the correlated tie-breaker version of uniform random serial dictatorship (U-RSD-T)), which is the baseline strategy-proof strict preferences baseline that is most commonly used in assignment markets. Note that accounting for indifferences increases the percentage of MBAs who get a first choice from 63% to 77%, a more than 20% improvement. Though indifferences are not the primary topic of this paper, it is striking that this sort of improvement can be made *without sacrificing strategy-proofness*. Indifferences are often treated as a theoretical inconvenience in the matching literature; the HBS provides evidence that failing to elicit indifferences can leave a significant amount of welfare on the table.




Figure 3: Results from the HBS match (no tie-breakers)

D Proofs

D.1 Rank efficiency

D.1.1 Characterization of rank-value mechanisms

Theorem 1. x is a rank efficient assignment $\Leftrightarrow \exists$ a valuation v such that x is supported by the v -rank-value mechanism.

Proof. We start by proving the first statement. The objective of the defining linear program can be rewritten as $\sum_{k=1}^{\bar{k}-1} N_{\alpha}^x(k) (v_k - v_{k+1})$. By way of contradiction, let \tilde{x} be a random assignment that α -rank-dominates the allocation yielded by the (α, v) -rank-efficient mechanism, x^* . Then, the following is true:

$$\sum_{k=1}^{\bar{k}-1} N_{\alpha}^{\tilde{x}}(k) \cdot (v_k - v_{k+1}) > \sum_{k=1}^{\bar{k}-1} N_{\alpha}^{x^*}(k) \cdot (v_k - v_{k+1})$$

because the rank-dominance of \tilde{x} over x^* means that the left-hand side is weakly greater than the right-hand side, term by term, and strictly so for at least one term. But this is a contradiction of the optimality of the rank-efficient mechanism's defining linear program.

Now, we prove the second statement. Consider the set of α -rank-CDFs that can be generated by a random allocations, $A_{\alpha} \subset \mathbb{R}^{\bar{k}-1}$, where the first index of the vector corresponds to $N_{\alpha}^x(1)$, the second to $N_{\alpha}^x(2)$, and so on. By the definition of an α -rank-CDF, A_{α} is convex. Now, let x be an α -rank-efficient allocation, and let $b \in A$ be the α -rank-CDF generated by x . Let the set of α -rank-CDFs that would α -rank-dominate the CDF b if they were generated by feasible random allocations be denoted by $U(b) \equiv \left\{ a \mid a \in \mathbb{R}^{\bar{k}-1}; a_i \geq b_i, \forall i; \exists i', a_{i'} > b_{i'} \right\}$. $U(b)$ is also convex. Any member of both A_{α} and $U(b)$ would α -rank-dominate x , so it must be that $A_{\alpha} \cap U(b) = \emptyset$. The separating hyperplane theorem then tells us that there is some $p \neq 0$ and some c such that $p \cdot u \geq c, \forall u \in U(b)$ and $p \cdot a \leq c, \forall a \in A_{\alpha}$. By construction, $p \cdot b = c$. Without loss of generality, assume $c \geq 0$. Now, let e_i denote the unit vector that points along the i^{th} axis of $\mathbb{R}^{\bar{k}-1}$. For any $\delta > 0$, we know that $b + \delta e_i \in U(b)$. This means that $p \cdot (b + \delta e_i) = c + \delta p_i \geq c$, which means that $p_i \geq 0$. If $p_i > 0, \forall i$, then we can rescale p and c such that $\sum_{j=1}^{\bar{k}-1} p_j = 1$. We will show this in the next paragraph,

but first, consider the valuation v where $v_i = 1 - \sum_{j=1}^{i-1} p_j, \forall i \in \{2, \dots, \bar{k} - 1\}, v_1 = 1$,

and $v_{\bar{k}} = 0$. Since $p_i > 0$, $v_k > v_{k+1}$, that is, we have indeed constructed a valuation. The (α, v) -rank-efficient mechanism supports x .

To show that $p_i > 0, \forall i$, we need the stronger Polyhedral Separating Hyperplane Theorem of McLennan (2002). Adopting the notation of that paper's Theorem 2, we let $P = A_\alpha - \mathbb{R}^{\bar{k}-1}$ and $b = p \in P$. The α -rank-efficiency of x tells us that x is not in the relative interior of P , so the Theorem lets us conclude that there is a hyperplane H such that P is contained in one of its half-spaces, $P \subset H^-$, and $F \equiv P \cap H$ is the smallest face of P that contains b . Let this hyperplane be $H = \{u \in \mathbb{R}^{\bar{k}-1} | p \cdot u = c\}$. First, note that for $\delta > 0$, $b - \delta e_i \in P$, by definition. Now, say that $p_i = 0$. Then, Lemma 2 of McLennan (2002) with $S = \{b\}$ tells us that the relative interior of S , that is b , is contained in the relative interior of the smallest face that contains S , that is F . Hence, for some $\delta > 0$ small enough, $b + \delta e_i \in F$. But since $F \subset P$, this contradicts our original assumption of the α -rank-efficiency of x . Hence, $p_i > 0, \forall i$. \square

D.1.2 Ex ante efficiency interpretation

The proof in the main text only addresses uniform weights. To justify an (α, v) -rank-value mechanism, we need $\int v_{a'k} \cdot F(v) = c_{a,a'} \cdot \int v_{ak} \cdot F(v)$. In other words, we need the unconditional expectations of v_a for all agents to be proportional.

D.1.3 Tough decisions interpretation

Lemma 1. *Without changing the number of copies of the objects (including \emptyset), consider adding $\sum_o q_o - |\mathcal{A}|$ dummy agents to the market that are indifferent between all the objects. Let \tilde{x} be an extension of x in which all real agents' allocations stay the same and the remaining objects are feasibly assigned to the dummy agents. Then, x is non-wasteful $\Leftrightarrow \nexists$ an improvement cycle in \tilde{x} that includes a dummy agent.*

Proof. (\Rightarrow) If there is an improvement cycle that includes a dummy agent, then there is also a two agent improvement cycle where one of the agents is the dummy. Since the dummy is indifferent between all the objects and is holding something that was not assigned in x , this cycle is equivalent to the real agent claiming something that was not assigned, contradicting the assumption that x was non-wasteful. (\Leftarrow) If x is wasteful, then \tilde{x} must have a two person improvement cycle between the real agent who wants to claim an object unassigned under x and the dummy agent who was assigned that object in \tilde{x} . \square

Corollary to the proof. *In the same setup, x is non-wasteful $\Leftrightarrow \nexists$ an (α, v) -rank improving cycle in \tilde{x} that includes a dummy agent, for any (α, v) .*

Proof. The same logic holds because a two-agent improvement cycle involving one real agent and one dummy is also an (α, v) -rank improving cycle for any (α, v) , and vice-versa. \square

Proposition 6. *x is ordinally efficient $\Leftrightarrow x$ is non-wasteful and admits no improvement cycles.*

Proof. This is a straightforward generalization of a theorem from Katta and Sethuraman (2006), which in turn is a generalization of a theorem from Bogomolnaia and Moulin (2001). The lemmas show us that we can generalize our setup to the setup of Katta and Sethuraman (2006), where $|\mathcal{A}| = \sum_o q_o$. Non-wastefulness bridges the gap. \square

Theorem 2. *x is rank efficient $\Leftrightarrow x$ is non-wasteful and \exists a valuation v such that x admits no v -rank-improving cycles.*

Proof. (\Rightarrow) If x is wasteful, then by Claim ?, it cannot be generated by any rank-value mechanism, and hence is not α -rank efficient by Theorem ?. Now, assume that for any valuation v , there exists an (α, v) -rank improving cycle. This would yield an improvement in the objective of the (α, v) -rank-value mechanism, which means that x cannot be generated by any (α, v) -rank-value mechanism, which, by Theorem ?, means that x is not α -rank efficient.

(\Leftarrow) Consider the augmented market of Lemma ?. Now, by way of contradiction, let \hat{x} α -rank dominate \tilde{x} in that market. Let $\xi \equiv \hat{x} - \tilde{x}$. Now, construct a cycle by the following procedure. Find some (a, o) such that $\xi_{ao} < 0$, and let that agent and object be (a_1, o_1) . Now, find some a such that $\xi_{ao_1} > 0$, and let that be a_2 . Necessarily there will be an o_2 such that $\xi_{a_2o_2} > 0$. Continue this process until we come back to (a_1, o_1) , which must happen by the finiteness of the market. By construction, we have a trade cycle in \tilde{x} . Now, define

$$\xi_{ao}^1 = \begin{cases} -\min_{(a', o') \in \tau} |\xi_{a'o'}| & \exists k \text{ s.t. } (a, o) = (a_k, o_k) \\ \min_{(a', o') \in \tau} |\xi_{a'o'}| & \exists k \text{ s.t. } (a, o) = (a_{k+1}, o_k) \\ 0 & \text{otherwise} \end{cases}$$

Then $\xi - \xi^1$ has at least one fewer non-zero entries than ξ . Hence, we can repeat this process to decompose ξ into a finite sum of trade cycles, that is $\xi = \sum_i \xi^i$ where each ξ^i represents a trade cycle. This means that $\hat{x} = \tilde{x} + \sum_i \xi^i$. Now, since \hat{x} α -rank

dominates, the objective of the (α, v) -rank-value mechanism is bigger at \hat{x} , regardless of the valuation. By the linearity of the objective of the (α, v) -rank-value mechanism, it cannot be that the value of the objective evaluated at ξ^i is weakly negative for all i . In other words, regardless of the valuation, at least one of the ξ^i 's must be (α, v) -rank improving, a contradiction. \square

D.2 Truth-telling in low information environments

We will derive all theorems allowing for indifferences, but we will state when an immediate corollary of the more general theorem is the less general theorem stated in the main text of the paper. Let $RVM_a^{\alpha, v}(\succsim; q; \tau)$ denote a 's allocation under the (α, v) -rank-value mechanism when the submitted preferences are \succsim , the capacities are q , and the tie-breaker ordering is τ .

Lemma 2. *Let $o' \succ_a o$. Then,*

$$[RVM_a^{\alpha, v}(\succsim; q; \tau) = o] \Rightarrow [RVM_a^{\alpha, v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; q; \tau) = o]$$

Proof. Let the objective function of the (α, v) -rank-value mechanism evaluated at assignment x and submitted preferences \succsim be denoted by

$$V^{\alpha, v}(\succsim; x) \equiv \sum_{a''} \sum_{o''} \alpha_{a''} \cdot v_{r_{a''}(o'')} \cdot x_{a''o''}$$

Then,

$$V^{\alpha, v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; x) = V^{\alpha, v}(\succsim; x) + \alpha_a \cdot (v_{r_a(o')} - v_{r_a(o)}) \cdot (x_{ao} - x_{ao'}) \quad (1)$$

Now, let $y = REM_a^{\alpha, v}(\succsim; q; \tau)$. Hence, $y_{ao} = 1$ and

$$V^{\alpha, v}(\succsim; y) - V^{\alpha, v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; y) = -\alpha_a \cdot (v_{r_a(o')} - v_{r_a(o)}) < 0$$

Now consider another assignment $z \neq y$ that is a member of the arg max of the linear program when when o and o' are switched, that is

$$V^{\alpha, v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; z) \geq V^{\alpha, v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; y) \quad (2)$$

First, we show that $z_{ao} = 1$. Assume otherwise. Then, by Equation 1, we derive

$$V^{\alpha,v}(\succsim; z) - V^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; z) \geq 0 > V^{\alpha,v}(\succsim; y) - V^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; y)$$

If we add this inequality to Inequality 2, we find that y was not an optimum to start with, a contradiction. Hence, a must receive o when he switches. Now, if $z_{ao} = 1$, then following the same logic, we derive

$$V^{\alpha,v}(\succsim; z) - V^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; z) = V^{\alpha,v}(\succsim; y) - V^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; y)$$

and combine it with Inequality 2 to find that either y was not an optimum to start with (a contradiction), or that z is a member of the arg max under both normal and switched preferences. But then, the only way that z could be chosen under the switch is if $\tau(z) < \tau(y)$, which would contradict y being chosen under the truth. \square

Corollary (to Lemma 2). *Let $o' \succ_a o$. Then,*

$$\left[RVM_a^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; \tau) = o' \right] \Rightarrow [REM_a^{\alpha,v}(\succsim; \tau) = o']$$

Proposition 7. *Under a rank-value mechanism, if agent a 's beliefs are $\{o, o'\}$ -symmetric, and $o' \succ_a o$, then for a , the allocation under any submitted preference that declares $o \succ_a o'$ is weakly stochastically dominated by the allocation of a submitted preference that does not. If the beliefs are symmetric, then this holds for all $o, o' \in \mathcal{O} \setminus \{\emptyset\}$.*

Proof. First, note that since the linear program does not depend on labels,

$$\begin{aligned} [RVM_a^{\alpha,v}(\succsim; q; \tau) = v] &\Leftrightarrow [RVM_a^{\alpha,v}(\succsim^{o \leftrightarrow o'}; q^{o \leftrightarrow o'}; \tau^{o \leftrightarrow o'}) = v] \\ [RVM_a^{\alpha,v}(\succsim; q; \tau) = o] &\Leftrightarrow [RVM_a^{\alpha,v}(\succsim^{o \leftrightarrow o'}; q^{o \leftrightarrow o'}; \tau^{o \leftrightarrow o'}) = o'] \end{aligned}$$

where $o, o' \in \mathcal{O}$ and $v \in \mathcal{O} \setminus \{o, o'\}$. Then, it must also be that

$$\begin{aligned} [RVM_a^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; q; \tau) = v] &\Leftrightarrow [RVM_a^{\alpha,v}(\succsim_a, \succsim_{-a}^{o \leftrightarrow o'}; q^{o \leftrightarrow o'}; \tau^{o \leftrightarrow o'}) = v] \\ [RVM_a^{\alpha,v}(\succsim_a^{o \leftrightarrow o'}, \succsim_{-a}; q; \tau) = o] &\Leftrightarrow [RVM_a^{\alpha,v}(\succsim_a, \succsim_{-a}^{o \leftrightarrow o'}; q^{o \leftrightarrow o'}; \tau^{o \leftrightarrow o'}) = o'] \end{aligned}$$

Now, using Lemma 2 and its corollary, we can think about the gains and losses an agent a who prefers o' to o gets from switching those two objects. The possible cases

		Lie: $RVM_a^{\alpha,v}(\succsim_a^{o\leftrightarrow o'}, \succsim_{-a}; q; \tau)$		
		$= u \notin \{o, o'\}$	$= o'$	$= o$
Truth: $RVM_a(\succsim_a, \succsim_{-a}; q; \tau)$	$= v \notin \{o, o'\}$	Case A	Impossible	Case B
	$= o'$	Case C	Case D	Case E
	$= o$	Impossible	Impossible	Case F

Table 4: Cases

	a tells the truth		a switches o and o'	
	$(\succsim_a, \succsim_{-a}; q; \tau)$	$(\succsim_a, \succsim_{-a}^{o\leftrightarrow o'}; q^{o\leftrightarrow o'}; \tau^{o\leftrightarrow o'})$	$(\succsim_a^{o\leftrightarrow o'}, \succsim_{-a}; q; \tau)$	$(\succsim_a^{o\leftrightarrow o'}, \succsim_{-a}^{o\leftrightarrow o'}; q^{o\leftrightarrow o'}; \tau^{o\leftrightarrow o'})$
Case A	v	u	u	v
Case B	v	o'	o	v
Case C	o'	u	u	o
Case D	o'	o	o'	o
Case E	o'	o'	o	o
Case F	o	o'	o	o'

Table 5: Assignments under the different cases

are shown in Table 4. Using these cases, we can then think about what our agent gets when everyone else's preferences switch o and o' , which is equally likely due to the assumption of $\{o, o'\}$ -symmetric beliefs. This analysis is shown in Table 5. In each of the cases, when we consider that everyone else switching o and o' is equally likely, truth-telling yields an allocation for a that weakly stochastically dominates what she would have gotten by switching o and o' . \square

We will formally name the condition in Proposition 7 weak order preservation. It essentially says that it is possible to break indifferences in such a way that the submitted preferences and the true preferences map to the same strict preference.

Definition 15. A reported preference \succsim'_a is called **weakly $\{o, o'\}$ -order-preserving with respect to \succsim_a** if $[o' \succ_a o] \Rightarrow [o' \succ'_a o]$. If this is true for all $o, o' \in \mathcal{O} \setminus \{\emptyset\}$, then we just call the reported preference **weakly order-preserving with respect to \succsim_a** .

Notice that if we are dealing with a world where only strict preferences are allowed, then this property basically says that objects besides \emptyset are listed in their true order, which is where we get Theorem 4. If we allow for indifferences, then we are not

insisting on truthful ordering, but instead are insisting that a true strict ordering of two objects is never reported reversed. Before moving on, we review a notation that was only briefly introduced in the main text. Let the true preference be \succsim_a and the preference that adds an indifference class I of unacceptable objects right above \emptyset be \succsim_a^{+I} . For some object o , let \succsim_a^{+o} denote $\succsim_a^{+\{o\}}$.

Lemma 3. *If the ranking scheme is upward-looking and unmatched-distinct then, submitting a preference where the bottom declared-acceptable indifference class consists entirely of truly unacceptable objects is weakly stochastically dominated by the submitting the same preference with the bottom declared-acceptable indifference class dropped.*

Proof. By way of contradiction, assume that submitting \succsim_a^{+I} makes a better off, that is $RVM_a(\succsim_a^{+I}, \succsim_{-a}; q; \tau) \equiv z$, $RVM_a(\succsim; q; \tau) \equiv y$, and $y_a \prec_a z_a$. Then, by optimality, we know that $V(\succsim; y) \geq V(\succsim; z)$ and $V(\succsim_a^{+I}, \succsim_{-a}; z) \geq V(\succsim_a^{+I}, \succsim_{-a}; y)$. Now, y_a is necessarily not in the lowest indifference class of \succsim_a^{+I} , but it could either be in a higher indifference class, or it could be \emptyset .

Case 1. (y_a is ranked above \succsim_a^{+I}) z_a must also be ranked above the bottom acceptable indifference class of \succsim_a^{+I} , so (since we have assumed that the rankings are upward-looking) $V(\succsim_a^{+I}, \succsim_{-a}; y) = V(\succsim; y)$ and $V(\succsim_a^{+I}, \succsim_{-a}; z) = V(\succsim; z)$. Combined with the optimality inequalities, these equalities imply that y and z are both in the arg max under both preferences. This cannot be, since the tiebreaker τ would choose one or the other.

Case 2. ($y_a = \emptyset$) z_a is necessarily not in the lowest acceptable indifference class of \succsim_a^{+I} , so $V(\succsim_a^{+I}, \succsim_{-a}; z) = V(\succsim; z)$. Now since the ranking scheme is unmatched-distinct, $V(\succsim_a^{+I}, \succsim_{-a}; y) = V(\succsim; y)$. Stringing inequalities together, we find that y and z are both in the arg max under both preferences. This cannot be, since the tiebreaker τ would choose one or the other.

□

Lemma 4. *If the ranking scheme is upward-looking and unmatched-neutral, then if o is a truly acceptable object that is declared unacceptable in \succsim_a , then*

$$RVM_a^{\alpha, v}(\succsim; q; \tau) \succsim_a RVM_a^{\alpha, v}(\succsim_a^{+o}, \succsim_{-a}; q; \tau)$$

Proof. Define $y \equiv RVM^{\alpha,v}(\succsim; q; \tau)$ and $z \equiv RVM^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; q; \tau)$. Since y is an arg max under \succsim , we have $V^{\alpha,v}(\succsim; y) \geq V^{\alpha,v}(\succsim; z)$, and since z is an arg max under $(\succsim_a^{+o}, \succsim_{-a})$, we have $V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; z) \geq V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; y)$.

Now, by way of contradiction, assume that $y_a \succ_a z_a$. There are three cases.

Case 1. ($z_a = \emptyset$) The ranking scheme is unmatched-neutral, so $V^{\alpha,v}(\succsim; z) \geq V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; z)$. Since y_a is above the part of the rank-order that has changed, we know $V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; y) = V^{\alpha,v}(\succsim; y)$ by the upward-looking assumption. Combined with the optimality inequalities, these equalities imply that y and z are both in the arg max under both preferences. This cannot be, since the tiebreaker τ would choose one or the other.

Case 2. ($z_a = o$) Denote by $z|^{(a,\emptyset)}$ the assignment we get by starting with z and only moving a from o to \emptyset . This is feasible, so since y is an arg max under \succsim , we know that $V^{\alpha,v}(\succsim; y) \geq V^{\alpha,v}(\succsim; z|^{(a,\emptyset)})$. By unmatched-neutrality, $V^{\alpha,v}(\succsim; z|^{(a,\emptyset)}) = V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; z)$. Again, since y_a is above any changes to the rank orders, $V^{\alpha,v}(\succsim; y) = V^{\alpha,v}(\succsim_a^{+o}, \succsim_{-a}; y)$. Combined with the fact that z is an arg max under $(\succsim_a^{+o}, \succsim_{-a})$, we can deduce that all of these objective values must be equal. But if this were true, there would be no way that the same tiebreaker τ could have chosen two different assignments.

Case 3. ($z_a \succ_a o$) Both z_a and y_a are above any changes in the rank orders. Hence, by the upward-looking assumption, z and y must be in the arg max for both submitted preferences, and since we break ties by choosing the lowest τ , it is a contradiction that we choose y under the true preferences and z under the truncations. □

Proposition 12. *Assume that agent a 's beliefs are symmetric. If the ranking scheme is upward-looking and unmatched-neutral, then truly acceptable objects must be declared acceptable. If the ranking scheme is upward-looking and unmatched-distinct, then any unacceptable objects that a declares acceptable must be in the bottom indifference class and declared indifferent to some acceptable objects.*

Proof. The previous lemma establishes the first part of the theorem. For the second part, note that symmetry requires that the objects are submitted in a weakly order preserving way. Hence, there can exist only one indifference class that contains both unacceptable and acceptable objects. Below that indifference class, all classes must be filled with unacceptable objects. By Lemma 3, we can drop those indifference classes. □

This combined with Proposition 7 gives us the generalization of Theorem 5 to preferences with indifferences.

Theorem 9. *If an agent's beliefs are symmetric and the rank scheme is upward-looking and...*

- *...unmatched-distinct, then his expected allocation from any other strategy is weakly stochastically dominated by his expected allocation from submitting some weakly order preserving preference in which any truly unacceptable objects that are declared acceptable are declared indifferent to some truly acceptable object and are declared in the lowest acceptable indifference class.*
- *...unmatched-neutral, then his expected allocation from any other strategy is weakly stochastically dominated by his expected allocation from submitting some weakly order preserving preference in which all truly acceptable objects are declared acceptable.*

D.3 Supportable rank efficiency and the competitive equilibrium mechanisms

Theorem 6. *Not all ordinally efficient allocations are supportably rank-efficient.*

Proof. For a trade cycle τ , define $\Delta_{\tau}^{(\alpha,v)}V = \sum_{m=1}^{|\tau|} \alpha_{i_m} \cdot (v_{r_{a_m}(o_{m-1})} - v_{r_{a_m}(o_m)})$, where o_0 is understood to be o_m . By Theorem ?, an assignment is only supportably rank efficient there exists some (α, v) such that $\Delta_{\tau}^{(\alpha,v)}V \leq 0$ for all trade cycles τ . Now, consider a 6 agent example, where there is one seat at a through d and 2 seats at e .

1 :	a		\succ	d	\succ	e
2 :	a	\succ	c		\succ	e
3 :	a	\succ	b		\succ	e
4 :		b	\succ	c		\succ e
5 :			c	\succ	d	\succ e
6 :	b				\succ	e

Probabilistic serial yields the random allocation

	a	b	c	d	e
1 :	1/3	0	0	17/27	1/27
2 :	1/3	0	7/27	0	11/27
3 :	1/3	1/9	0	0	15/27
4 :	0	4/9	4/27	0	11/27
5 :	0	0	16/27	10/27	1/27
6 :	0	4/9	0	0	15/27

which we know to be ordinally efficient. Relative to the trade cycles

$$\begin{aligned}
\tau_1 &= ((1, a), (2, e)) \\
\tau_2 &= ((1, e), (2, a)) \\
\tau_3 &= ((1, a), (3, b), (4, c), (5, d)) \\
\tau_4 &= ((1, a), (2, e), (3, b), (4, c), (5, d))
\end{aligned}$$

we can calculate how $\Delta_{\tau}^{(\alpha, v)}V$ will change:

$$\begin{aligned}
\Delta_{\tau_1}^{(\alpha, v)}V &= (\alpha_2 - \alpha_1) \cdot (v_1 - v_3) \\
\Delta_{\tau_2}^{(\alpha, v)}V &= (\alpha_1 - \alpha_2) \cdot (v_1 - v_3) \\
\Delta_{\tau_3}^{(\alpha, v)}V &= (\alpha_3 + \alpha_4 + \alpha_5 - \alpha_1) \cdot (v_1 - v_2) \\
\Delta_{\tau_4}^{(\alpha, v)}V &= -\alpha_1 \cdot (v_1 - v_2) + \alpha_2 \cdot (v_1 - v_3) - \alpha_3 \cdot (v_2 - v_3) + (\alpha_4 + \alpha_5) \cdot (v_1 - v_2)
\end{aligned}$$

Now, we show how there is no (α, v) such that one of these trade cycles isn't (α, v) -rank improving, that is $\Delta_{\tau}^{(\alpha, v)}V \leq 0$. Since $v_1 > v_3$, we need $\alpha_1 = \alpha_2$ to prevent either τ_1 and τ_2 from being (α, v) -rank-improving. To prevent τ_3 from being α -rank-improving, we need $\alpha_1 \geq \alpha_3 + \alpha_4 + \alpha_5$. Finally, since $\alpha_1 = \alpha_2$, we can simplify $\Delta_{\tau_4}^{(\alpha, v)}V$ to $(\alpha_1 - \alpha_3) \cdot (v_2 - v_3) + (\alpha_4 + \alpha_5) \cdot (v_1 - v_2)$. The second term is definitely positive, so it must be that $\alpha_1 < \alpha_3$. But combining the last two inequalities gives us $\alpha_3 > \alpha_3 + \alpha_4 + \alpha_5$, which cannot be. Hence, there is no (α, v) such that $\Delta_{\tau}^{(\alpha, v)}V \leq 0$ for all four of the trade cycles we have listed. Thus the ordinally efficient assignment we started with is not supportably rank efficient.

□

Proposition 8. x is α -rank efficient $\Leftrightarrow \exists \hat{p}, v$ such that (x^*, \hat{p}) form an (α, v) -discount equilibrium

Proof. \Rightarrow

The LP dual of the (α, v) -rank efficient mechanism is

$$\begin{aligned} \min_{u_a \geq 0, p_o \geq 0} & \left\{ \sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot u_a \right\} \\ \text{s.t.} & \quad u_a \geq v_{r_a(o)} - \frac{p_o}{\alpha_a} \quad \forall o \in \mathcal{O}, a \in \mathcal{A} \end{aligned}$$

First, note that some prices must be zero. To see this, say that they weren't. Then we could subtract some $\varepsilon > 0$ from all prices and add ε/α_a to u_a . This change does not violate any constraints and adds the term $(|\mathcal{A}| - \sum_o q_o) \cdot \varepsilon$ to the objective. Since we are assuming that $\emptyset \in \mathcal{O}$, this term must be negative, which contradicts optimality.

We will support the discount equilibrium with the prices from the dual of the rank efficient mechanism. Now, the LP dual of the agent a 's problem in the (α, v) -discount equilibrium with these prices is

$$\begin{aligned} \min_{u_a} & \quad u_a \\ \text{s.t.} & \quad u_a \geq v_{r_a(o)} - \frac{p_o}{\alpha_a} \quad \forall o \in \mathcal{O} \end{aligned}$$

By nested optimization, we see that the u_a from the rank efficient mechanism's dual also solves the discount equilibrium duals. Now, let x_{ao} be the assignment we are considering from the arg max of the rank efficient mechanism. The LP duality theorem tells us that

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot u_a = \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao}$$

From the constraint in the dual of the rank efficient mechanism, we can then derive that, for any ξ such that $\sum_o \xi_{ao} \leq 1, \forall a \in \mathcal{A}$, we have

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_{o \in \mathcal{O}} \left(v_{r_a(o)} - \frac{p_o}{\alpha_a} \right) \cdot \xi_{ao} \leq \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao}$$

If we plug in x for ξ , then this reduces to $\sum_{o \in \mathcal{O}} p_o \cdot (q_o - \sum_{a \in \mathcal{A}} x_{ao}) \leq 0$. Since all elements of that sum are weakly positive, we conclude that $\sum_{o \in \mathcal{O}} p_o \cdot (q_o - \sum_{a \in \mathcal{A}} x_{ao}) = 0$. This, coupled with a constraint from the rank efficient mechanism, is the second condition of the discount equilibrium. Note that we have also shown that

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_{o \in \mathcal{O}} \left(v_{r_a(o)} - \frac{p_o}{\alpha_a} \right) \cdot x_{ao} = \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao}$$

Finally, we show that the x_{ao} from the arg max of the rank efficient mechanism must also be part of the arg max for the agents' problems in the discount equilibrium. By way of contradiction, say that they aren't, and let ξ_a be the bundle chosen by agent a under prices p . Then, for all agents, $\sum_o \xi_{ao} \leq 1$ and $\sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot \xi_{ao} \geq \sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot x_{ao}$, with the second inequality strict for some agent a' . Thus,

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot \xi_{ao} > \sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot x_{ao}$$

But we have shown

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot x_{ao} = \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao},$$

which means that we have derived

$$\sum_{o \in \mathcal{O}} q_o \cdot p_o + \sum_{a \in \mathcal{A}} \alpha_a \cdot \sum_o \left(v_{r_a(o)} - \frac{p_o}{\alpha_a}\right) \cdot \xi_{ao} > \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} \alpha_a \cdot v_{r_a(o)} \cdot x_{ao}$$

which stands in direct contradiction to something we have already proven. Therefore, the assignment from the rank efficient mechanism solves the agents' problems in the discount equilibrium. Thus we have shown that any member of the arg max of the rank-efficient equilibrium can be supported as a discount equilibrium.

◀

Consider the optimization

$$\begin{aligned} \max_x \quad & \left\{ \sum_o \sum_a \alpha_a \cdot v_{r_a(o)} \cdot x_{ao} + \sum_o p_o \cdot \left(q_o - \sum_a x_{ao} \right) \right\} \\ \text{s.t.} \quad & \sum_o x_{ao} \leq 1 \\ & \sum_a x_{ao} \leq q_a \end{aligned}$$

Clearly the value of this optimization is less than or equal to the value of

$$\begin{aligned} \max_x \quad & \sum_o \sum_a \alpha_a \cdot v_{r_a(o)} \cdot x_{ao} \\ \text{s.t.} \quad & \sum_o x_{ao} \leq 1 \\ & \sum_a x_{ao} \leq q_a \end{aligned}$$

Now, if we start with the assignment x from some discount equilibrium, by nested

optimization and market clearing, we know that it solves the first optimization. But by the second condition in the definition of a discount equilibrium, we know that the value of this assignment in the first optimization is attained in the second as well. Hence, x must solve the second optimization, since it attains the upper bound at x . Hence, x is in the arg max of the rank efficient mechanism's defining program. \square

Theorem 7. x is supportably rank efficient $\Leftrightarrow \exists(B, v)$ such that x is supported by the (B, v) -budget mechanism.

Proof. We have already proven that discount equilibria and α -rank efficient assignments are the same things. Now, we show that discount equilibria and budget equilibria are the same. This established the theorem.

(DE \Rightarrow BE) Take an (α, v) -discount equilibrium, (x^*, \hat{p}) . Now, let $B_a = \sum_o \hat{p}_o \cdot x_{ao}^*$ be the income assigned to agent a . Now, by way of contradiction, assume that (x^*, \hat{p}) is not a (B, v) -budget equilibrium assignment for B . By construction, it is feasible, so it must be that there is some other feasible \tilde{x} such that $\sum_o v_{r_a(o)} \cdot \tilde{x}_{ao} > \sum_o v_{r_a(o)} \cdot x_{ao}^*$. Since it is feasible, we also know that $\sum_o \hat{p}_o \cdot \tilde{x}_{ao} \leq \sum_o \hat{p}_o \cdot x_{ao}^*$. Combining the last two inequalities (multiplying the first by α_a), we find $\sum_o (v_{r_a(o)} - \hat{p}_o) \cdot \tilde{x}_{ao} > \sum_o (v_{r_a(o)} - \hat{p}_o) \cdot x_{ao}^*$, a contradiction of the original assumption that (x^*, \hat{p}) was an (α, v) -assignment equilibrium.

(BE \Rightarrow DE) Take a (B, v) -budget equilibrium, (\tilde{x}, \hat{p}) . Now, consider the LP dual of the agent optimization problem:

$$\begin{aligned} (\tilde{\lambda}_a, \tilde{\mu}_a) \in \arg \min_{\lambda_a, \mu_a \geq 0} \quad & \mu_a + \lambda_a \cdot \widehat{B}_a \\ \text{s.t.} \quad & \mu_a \geq v_{r_a(o)} - \hat{p}_o \cdot \lambda_a \end{aligned}, \forall a$$

Take a solution to this dual problem for an agent a . Also, consider the LP dual of the agent optimization problem in an (α, v) -discount equilibrium:

$$\begin{aligned} \mu_a^* \in \arg \min_{\mu_a \geq 0} \quad & \mu_a \\ \text{s.t.} \quad & \mu_a \geq \alpha_a \cdot v_{r_a(o)} - \hat{p}_o \end{aligned}, \forall a$$

Now, from here we continue by cases.

Case 1. ($\tilde{\lambda}_a = \mathbf{0}$): By inspection of the dual, we can see that $\tilde{\mu}_a = v_1$. By the LP duality theorem, we then know that $\sum_o v_{r_a(o)} \cdot \tilde{x}_{ao} = \tilde{\mu}_a = v_1$. The only way this can happen is if $\tilde{x}_{ao} = \mathbf{1}_{\{r_a(o)=1\}}$, that is, if agent a is able to buy a full share of his

most preferred object, $r_a^{-1}(1)$. without exhausting his budget. Set $\alpha_a = 1 + \frac{\widehat{p}_{r_a^{-1}(1)}}{v_1}$. This does two things. First, it makes the objective in the agent's optimization the same in both the discount and budget equilibrium (remember, only one term in the sum is actually there, since x_{ao} is only positive for one o). Hence, the agent's problem in the discount equilibrium is a relaxation of the agent's problem in the budget equilibrium. Second, it makes the LP duals of the optimizations in the two problems identical. These two facts, coupled with the LP duality theorem, mean that $\sum_o v_{r_a(o)} \cdot x_{ao}^* = \mu_a^* = \widetilde{\mu}_a = \sum_o v_{r_a(o)} \cdot \widetilde{x}_{ao}$. Since \widetilde{x} is feasible and attains the optimum, we have shown that our agent has found an optimum of the agent's problem in the discount equilibrium.

Case 2. ($\widetilde{\lambda}_a > \mathbf{0}$): Set $\alpha_a = \frac{1}{\widetilde{\lambda}_a}$. This makes our two LP duals rescaled versions of each other, such that $\mu_a^* = \frac{\widetilde{\mu}_a}{\widetilde{\lambda}_a} + B_a$. The LP duality theorem then gives us that $\sum_o (\alpha_a \cdot v_{r_a(o)} - \widehat{p}_o) \cdot x_{ao}^* = \mu_a^* = \frac{\widetilde{\mu}_a}{\widetilde{\lambda}_a} = \sum_o \alpha_a \cdot v_{r_a(o)} \cdot \widetilde{x}_{ao} - B_a$. This means that $\sum_o \widehat{p}_o \cdot x_{ao}^* = B_a$, that is x^* is feasible in the budget equilibrium agent optimization. Further, this means that $\sum_o \widehat{p}_o \cdot \widetilde{x}_{ao} \leq \sum_o \widehat{p}_o \cdot x_{ao}^*$, since a budget equilibrium assignment must also be feasible. Now, assume that the budget equilibrium assignment is not part of a discount equilibrium. This must mean that $\sum_o (\alpha_a \cdot v_{r_a(o)} - \widehat{p}_o) \cdot x_{ao}^* > \sum_o (\alpha_a \cdot v_{r_a(o)} - \widehat{p}_o) \cdot \widetilde{x}_{ao}$. Combining the last two inequalities and factoring out a common α_a gives us $\sum_o v_{r_a(o)} \cdot x_{ao}^* > \sum_o v_{r_a(o)} \cdot \widetilde{x}_{ao}$ which contradicts the original assertion that \widetilde{x}_{ao} was part of a discount equilibrium assignment. Hence, \widetilde{x} is an optimum of the agent's problem in the discount equilibrium.

Hence we have shown that agents optimizing according to the problem in the discount equilibrium can choose \widetilde{x} , given an appropriately chosen α vector. \square

The efficiency of the budget equilibria lines up well with our understanding of general equilibrium and the welfare theorems, but why should the discount equilibrium yield an efficient outcome. Distorting prices with the α -discounts seems like it should cause trouble, in the same way that distortionary taxation does. So what is going on? We can think of the agent's problem in the discount equilibrium as an agent with enough money to buy whatever he wants, but who values money quasi-linearly. Since we have this sort of separation, and our efficiency concepts only rely on the distribution of objects among the agents, the distortion only serves to limit how much each agent is willing to spend on objects. In short, in the discount equilibrium, distortions in the relative cost of keeping money serve a similar role to the budgets in the budget equilibrium.