

Rank Efficiency: Modeling a Common Policymaker Objective[†]

Clayton R. Featherstone

The Wharton School, University of Pennsylvania

NOVEMBER 10, 2020

Abstract

Policymakers often gauge matches using *rank distributions*—how many get their reported 1st choice, 2nd choice, etc. Formalizing this observation, an assignment is *rank efficient* if its rank distribution cannot be stochastically dominated. Rank efficiency refines ordinal efficiency, and hence ex post (Pareto) efficiency. In addition, a class of linear-programming mechanisms seen in the field guarantee rank-efficient assignments. Policymakers can also attain rank efficiency simply by looking for local improvements that increase a natural objective. In a Harvard Business School match, such *tinkering* could have increased the number who get their first or second choice by 18 percent. Such gains suggest tinkering might be widespread, which magnifies rank efficiency’s importance as a descriptive concept. However, since rank efficiency and strategyproofness cannot coexist, tinkering gains may prove illusory. Nevertheless, tinkering need not undermine incentive properties entirely: when agents have little information and no outside options, truth-telling can be a best-response.

ACKNOWLEDGMENTS: I am grateful to many people for helpful comments on this paper, including (but not limited to) Al Roth, Muriel Niederle, Jacob Leshno, Judd Kessler, Corinne Low, Juuso Toikka, and John Hatfield.

[†]. This paper has also circulated under the title “Rank Efficiency: Investigating a Widespread Ordinal Welfare Criterion.”

1 Introduction

INSTITUTIONS that use ordinal assignment mechanisms often gauge success by looking at **rank distributions**, that is, how many participants get their reported first choice, second choice, and so on. For example, when school districts announce the results of a choice-based placement system, they often report things like how many students were assigned to their top choice, one of their top three choices, or one of their ranked choices at all. (See New York City Department of Education 2009, 2018 and San Francisco Unified School District 2011, 2018a for publicly released reports that demonstrate this.) Sometimes organizations use rank distributions internally to evaluate new mechanisms, as the San Francisco Unified School District (SFUSD) did in 2010, and as Harvard Business School (HBS) did in designing the process used to match MBAs to the country in which they fulfill the Field Immersion Experiences for Leadership Development (FIELD) requirement.¹ And in fact, some institutions even pick an assignment explicitly based on rank distribution, as Teach For America (TFA) does when it assigns teachers to regions.² These examples aren't terribly surprising—after all, the rank distribution is a natural summary statistic for the quality of a match. However, the market design literature has tended to focus on more standard economic concepts like strategyproofness and Pareto efficiency.

The main purpose of this paper is to formally model the use of rank distribution-based objectives. By doing so, I construct a descriptive model of the assignments and mechanisms to which policymakers naturally gravitate. Key to my approach is the concept of **rank efficiency**. An overall assignment is rank efficient if its rank distribution cannot be first-order stochastically dominated by that of another assignment.

The paper is divided into two parts. The first (Sections 2–4) introduces the model, places rank efficiency in the context of other efficiency concepts, and characterizes rank-efficient mechanisms. The net result is a unifying framework for efficiency concepts in an ordinal setting. The second part of the paper (Sections 5 and 6) considers two practical issues: the strong temptation for policymakers to “tinker” with assignments to improve their rank distributions and the incentive for clearinghouse participants to deviate from truth-telling under a rank-efficient mechanism. Concerning the first issue, I show that an arbitrary mechanism can be effectively transformed into a rank-efficient mechanism through an ex post tinkering process where each change to the assignment locally improves a simple objective. If policymakers care deeply about rank distributions, this suggests that

-
1. Along with Atila Abdulkadiroğlu, Muriel Niederle, Parag Pathak, and Al Roth, I helped the SFUSD consider new assignment mechanisms over the course of the 2009–2010 school year. Al Roth and I helped to design the HBS FIELD match between 2010 and 2012.
 2. Teach For America is a nationwide nonprofit that sends mostly new college graduates to teach in low-performing public and charter schools around the US. Al Roth and I helped the organization to streamline its assignment process between 2010 and 2012.

tinkering might be widespread, which magnifies the importance of rank efficiency as a descriptive concept. Concerning the second practical issue, I show that, while no rank-efficient mechanism is strategyproof, truth-telling can be a best-response when participants have little information (à la Roth and Rothblum 1999) and no outside options. Hence, while rank-efficient mechanisms will be undermined by poor incentive properties in some settings, they might be salvageable in others. Throughout the paper, I will interpret results through the lens of my personal experience in helping to design the TFA and HBS FIELD matches (see footnotes 1 and 2).

1.1 Efficiency Concepts and Mechanisms (Sections 2–4)

The first part of the paper begins by showing that rank efficiency is a refinement of the other main one-sided efficiency concepts—ex post (Pareto) efficiency and ordinal efficiency (Bogomolnaia and Moulin 2001). This continues to be true, even on the domain of deterministic assignments, where by contrast, ordinal efficiency and ex post efficiency are equivalent. Intuitively then, to a policymaker who cares only about the realized deterministic assignment, ordinal efficiency doesn’t offer anything new over ex post efficiency, while rank efficiency does. Rank efficiency also captures the intuition that policymakers are sometimes dissatisfied with *mere* ex post efficiency. Pareto concepts never insist on hurting one participant to help many, while rank efficiency often does so.

An example clarifies the point. Let there be four objects (labeled o_1 through o_4) and four agents (labeled a_1 through a_4), with the preferences listed in Figure 1. Any object not listed in an agent’s preference is unacceptable to her. Now, compare the assignment where each agent gets the boxed object in her preference to the assignment where each agent gets the circled object. These two assignments are not Pareto comparable, yet the circled assignment gives three agents a 1st choice, while the boxed only gives one agent a 1st choice. The rank distributions are clearly stochastically ranked. Insisting that the only way to improve an assignment is to make all agents better off—which I call **Pareto agnosticism**—seems to miss a clear intuition. Most policymakers would choose the circled assignment, making the “tough decision” to hurt one agent to help the others.

Of course, automating this “tough decision” could be a problem if the policymaker wants to avoid hurting agent a_1 —she could, for instance, be a member of a socially disadvantaged group. That policymakers often report separate rank distributions for different classes of participants suggests this is not an idle concern. For example, the SFUSD reports the rank distribution by ethnicity and zipcode (San Francisco Unified School District 2018b), and the National Residency Matching Program (NRMP), which matches new physicians to American residency programs (Roth 1984; Roth and Peranson 1999; Kojima, Pathak, and Roth 2013), reports match rates by whether the graduate’s medical school is allopathic or osteopathic, and by whether it is

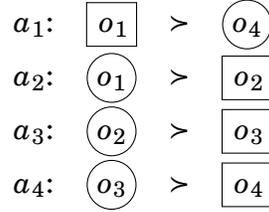


Figure 1: Two Assignments That Are Not Pareto Comparable

NOTES: Agents a_1 through a_4 have preferences over objects o_1 through o_4 . Any objects not listed in an agent’s preference are unacceptable to that agent. The circled (boxed) objects represent the assignment where each agent gets the circled (boxed) object in her preference.

foreign or domestic (National Residency Matching Program 2019).

To model such behavior, I generalize rank efficiency by partitioning the set of match participants and only requiring that there be no other assignment where *each partition element’s* rank distribution is stochastically better. This **partitioned rank efficiency** concept can insist on hurting one participant to help many, but only if those participants are all in the same partition element. I then show that partitioned rank efficiency bridges the gap between ordinal efficiency and rank efficiency: the former corresponds to the finest partition of the agents, while the latter corresponds to the coarsest.³

Next, I introduce the **welfare maximization mechanism**. For each participant, the mechanism assumes she places some cardinal utility on getting her first choice, some lower cardinal utility on getting her second choice, and so on. It then finds an assignment that maximizes the welfare sum, which can be done via a simple linear program. Similar mechanisms have been seen in the field, such as the match of house officers (the UK equivalent of US medical residents) to hospitals in Cambridge and London (Roth 1991; Ünver 2001, 2005) and the TFA match mentioned above.

When cardinal utility assumptions are unrestricted (i.e., allowed to vary freely across agents), the welfare maximization mechanism yields assignments that are only ordinally efficient. However, if participants are assumed to have the same cardinal utility profile (that is, for a given rank, k , everyone places the same cardinal value on getting a k th-choice allocation) then the welfare maximization mechanism yields assignments that are rank efficient. Similarly, given a partition, if each participant in a partition element is assumed to have the same cardinal utility profile, then the welfare maximization mechanism yields assignments that are partitioned rank efficient.

These mechanisms provide intuition for the way in which their underlying efficiency concepts deviate from Pareto agnosticism. Rank efficiency assumes that a k th-choice allocation is of the same social value, regardless

3. In addition, the partitioned rank efficiency concepts are naturally ordered: if one partition is a refinement of another, then any assignment that is partitioned rank efficient relative to the coarser partition must also be partitioned rank efficient relative to the finer partition.

of who gets it, while ordinal efficiency makes no such assumption. Partitioned rank efficiency lies between. Hence, the cardinal utility assumptions that feed into the welfare maximization mechanism provide a framework for policymakers to provide instructions on how to choose their preferred Pareto-efficient assignment. These assumptions also provide a framework that ties together many efficiency concepts found in the matching literature.

That said, making a distinct assumption about each agent’s cardinal utility profile is daunting, especially if policymakers lack information to distinguish between agents. A natural assumption—and the one codified by aggregating assignments into overall rank distributions—is that all participants have the same cardinal utility profile. I call a welfare maximization mechanism that makes this assumption (and hence yields rank-efficient assignments) a **rank-value mechanism**. Such a mechanism is straightforward: score assignments by giving some number of points for each 1st-choice allocation, some smaller number of points for every 2nd-choice allocation, and so on, and then pick the assignment with the highest score.

1.2 Tinkering and Incentives ([Sections 5 and 6](#))

The second part of the paper shifts to more practical concerns. It begins by looking at a phenomenon I call **tinkering**—changing an assignment ex post to improve the rank distribution. When I first encountered TFA’s match procedure, the organization got its initial assignment by running the teachers’ rankings through a computer algorithm of unknown provenance. TFA administrators felt this assignment had an inferior rank distribution. To remedy the situation, they spent about a week improving it by hand. This process involved looking for potentially non-Pareto-improving trades where the good was perceived to outweigh the bad. Every thirty minutes or so, progress was gauged by computing the rank distribution. The process continued until the rank distribution could no longer be improved.⁴

In [Section 5](#), I formally model this process with a simple algorithm: given an assumption about the values of differently ranked allocations, implement any (potentially non-Pareto-improving) trade that will increase the objective of the corresponding rank-value mechanism, and continue doing so until there are no more such trades. The algorithm is not meant as a mechanism to recommend to policymakers, but rather a model of how some policymakers might act behind closed doors. I show the algorithm is guaranteed to terminate at a rank-efficient assignment, which practically means that policymakers need not understand linear programming to implement a rank-value mechanism. It also means that any matching process followed by sufficient tinkering effectively becomes a rank-value mechanism.

Given the intuitive appeal of tinkering, one might suspect that it is

4. More recently, Bronfman et al. ([2015](#)) reports a similar tinkering procedure has been adopted by the Israeli Ministry of Health to match newly minted physicians to internships in that country.

widespread. However, this might not be the case if the return from tinkering isn't large. I empirically investigate by using ranking data from the 900 MBAs that participated in the HBS FIELD match. Since the mechanism is (essentially) random serial dictatorship, which is strategyproof (Roth and Sotomayor 1992), it is reasonable to treat the rankings as truthful (cf. Li 2017). By comparing the rank distributions generated by random serial dictatorship to those generated by a counterfactual rank-value mechanism, I show that tinkering increases the expected number of second-or-better allocations by 18 percent. In this light, it is unsurprising that in the first year the mechanism was run, HBS administrators seriously considered tinkering. More broadly, this suggests that when market designers implement a strategyproof mechanism, they should be aware how extreme the temptation for policymakers to tinker might be.

Of course, the gains from tinkering might be illusory if tinkering gives agents incentive to deviate from truth-telling. In Section 6, I show that no rank-efficient mechanism is strategyproof.⁵ This result had practical implications for HBS administrators: ultimately, they didn't tinker because students had been promised a mechanism that wouldn't penalize them for truthfully ranking countries. In fact, when choosing a match procedure, HBS considered and rejected a rank-value mechanism due to its lack of strategyproofness. Given previous experience at HBS (Budish and Cantillon 2012) and Wharton (Budish and Kessler 2018), it seemed reasonable to worry that the MBAs would quickly learn to game the system.

This concern about lack of strategyproofness was not shared by TFA: in the spring of 2011, it institutionalized and automated its tinkering approach by implementing a rank-value mechanism via linear programming. Administrators reasoned that teachers only apply once, are geographically separated, and know little about what regions are popular relative to capacity. I consider such reasoning by modeling participant beliefs as relatively uninformed, à la Roth and Rothblum (1999). Under such beliefs, I find some support for TFA's rationale: truth-telling is a best-response if participants are forced to rank all options, as they are in TFA's match.

That said, when a participant is allowed to express which options are acceptable (i.e., preferred to the outside option), she can no longer best-respond by truth-telling, even in the low-information environment alluded to in the previous paragraph. In fact, the form of her best-response depends on a seemingly insignificant design detail: how the mechanism treats leaving someone unmatched. Consider an agent who submits the ranking $o_1 > o_2 > \emptyset$, where o_1 and o_2 are genuine objects to be assigned and \emptyset represents being unmatched (i.e., the agent's outside option). Giving \emptyset to this agent could be treated as giving her an outcome that is, in some way, distinct from any other outcome, but it could also be treated as giving her a

5. Unlike the Bogomolnaia and Moulin (2001) impossibility result for ordinally efficient mechanisms, I am able to dispense with the *equal treatment of equals* assumption.

3rd choice. Clearinghouses in the field tend to report the number of unmatched agents as a distinct category (see the reports referenced in the first paragraph of this introduction), which resonates with the first method; however, the second method is quite sensible, as it seems better to leave a participant unmatched than to give her an assignment that she will refuse *ex post*.

If the mechanism’s treatment of unmatched agents follows the first method, then a participant with uninformed beliefs best-responds in a familiar way, by submitting a *truncation* of her true preferences (i.e., truthful except for possibly declaring some truly acceptable outcomes unacceptable). However, if the mechanism follows the second method, then a participant with uninformed beliefs has a novel and unique best-response: she should submit the *full extension* of her true preference (i.e., truthful, except the outside option is ranked last).

The rest of the paper is organized as follows. [Section 2](#) lays out the baseline model. [Section 3](#) defines rank efficiency and relates it to *ex post* efficiency and ordinal efficiency. [Section 4](#) introduces the welfare maximization mechanism and shows that it yields rank-efficient assignments under the proper assumptions. [Section 5](#) models the tinkering process and shows that tinkering has real gains in the context of the HBS match. This demonstrates the temptation that can be faced by policymakers when market designers implement a strategyproof mechanism and suggests the prevalence of tinkering more widely. [Section 6](#) investigates when policymakers can expect truth-telling to a rank-value mechanism and how agents respond to a lack of information more broadly. [Section 7](#) concludes. Proofs to all propositions can be found in [Appendix A](#), while a fuller description of the empirical exercise described in [Section 5](#) can be found in [Appendix B](#).

2 Model

EACH **agent** from the set \mathcal{A} must ultimately be allocated a copy of an **object** whose **type** is in the set \mathcal{O} . Going forward, when the variables o and a are mentioned without explicit ranges, those ranges should be understood as \mathcal{O} and \mathcal{A} , respectively. For each type, o , there is a positive integer number of copies, q_o , which I will call its **capacity**. The **capacity profile** is the vector $q = (q_o)$. I assume that there are enough objects for every agent to be matched, that is, $\sum_o q_o \geq |\mathcal{A}|$. This does not rule out agents remaining unassigned, but instead requires lack of assignment to be codified by a special **null** object type, \emptyset , that has sufficient capacity for all agents to be assigned a copy (i.e., that has $q_\emptyset = |\mathcal{A}|$).⁶ Each agent interprets the null object as her outside option. Hence, the presence of \emptyset in \mathcal{O} means it is valid to allocate an agent’s outside option to her; its absence means it is not.

6. I will never use \emptyset to represent the empty set; it always represents the null object type.

Each agent, a , has a strict **ranking** over the object types, which I will represent as an $|\mathcal{O}|$ -dimensional vector, r_a , where $r_{ao} = k$ means that object type o is ranked k th by agent a .⁷ When the null object type is present, object types ranked above it are called **acceptable**, while object types ranked below it are called **unacceptable**. The **profile of rankings** can be represented by the matrix $r = (r_{ao})$. An **ordinal market** is an ordered quadruple, $(\mathcal{A}, \mathcal{O}, r, q)$. I assume agents to be indifferent over the objects received by others.

Randomness in allocation is sometimes desirable, especially when dealing with fairness concerns like the equal treatment of equals (Thomson 2011). Hence, I allow agents to be allocated lotteries over object types, so long as all of these lotteries can be simultaneously resolved. I will return to this issue momentarily, but for now, consider a matrix, $x = (x_{ao})$, whose (a, o) entry represents the probability that agent a is allocated a copy of object type o . I call x an **assignment** if it is **feasible**, i.e., for all agents a and object types o , the entry x_{ao} is non-negative and the constraints $\sum_{o'} x_{ao'} = 1$ and $\sum_{a'} x_{a'o} \leq q_o$ hold. Non-negativity and the equality constrain each agent's allocation to be a valid lottery over object types, while the inequality ensures that the assignment doesn't give out more copies of each object type than there are. I will refer to the lottery assigned to agent a —given by the a th row of an assignment—as her **allocation**.

An assignment is **deterministic** if all of its entries are either zero or one. Coming back to the issue from the previous paragraph, all agents' allocations in an assignment can be simultaneously resolved if that assignment is the expectation of a random matrix for which all possible realizations are deterministic assignments. Fortunately, this is true of any assignment that meets the definition above (Budish et al. 2013; Birkhoff 1946). When the expectation of a random matrix of deterministic support equals some assignment x , I will call that random matrix an **implementation** of x .⁸ Note that an assignment can have multiple implementations with different supports; however, since all of them induce the same allocation for each agent, the issue is generally unimportant (cf. footnote 9).

Finally, I call a function that maps ordinal markets to assignments an **ordinal assignment mechanism**. To be clear, in this paper, each agent makes decisions according to her cardinal utility, but only submits ordinal information (i.e., her ranking) to the assignment mechanism.

7. At points in the exposition that follows, I will assume that each agent, a , also has an underlying $|\mathcal{O}|$ -dimensional **cardinal utility** vector, u_a , that **rationalizes** her ranking. That is, for any two object types, o and o' , the inequality $u_{ao} > u_{ao'}$ holds if and only if the inequality $r_{ao} < r_{ao'}$ holds. This cardinal utility vector represents a von Neumann–Morgenstern expected utility on lotteries over object types.

8. Note that my use of the term *implementation* follows Budish et al. (2013). It should not be confused for the usage in mechanism design, where a mechanism is said to implement a social-choice correspondence via some solution concept (see, e.g., Jackson 2001).

3 Efficiency Concepts

IN THIS section I begin by placing rank efficiency in the context of two other ordinal concepts: the standard, ex post efficiency, and its most well-known refinement, ordinal efficiency (introduced by Bogomolnaia and Moulin 2001, but see also McLennan 2002, Abdulkadiroğlu and Sönmez 2003, Manea 2008, Manea 2009, Che and Kojima 2010, and Kojima and Manea 2010). While both ordinal efficiency and rank efficiency are refinements of ex post efficiency (as I will show), two things differentiate rank efficiency from ordinal efficiency.

First, on the limited domain of deterministic assignments, I show that ordinal efficiency is equivalent to ex post efficiency, while rank efficiency remains a refinement of it. In other words, to a policymaker who cares only about the realized deterministic assignment, ordinal efficiency doesn't offer anything new over ex post efficiency, while rank efficiency does.

Second, unlike ex post and ordinal efficiency, rank efficiency is not Pareto agnostic (as defined in the introduction). By selecting a subset of the ordinally efficient assignments, rank efficiency makes “tough decisions” that hurt some agents to help others. Policymakers, by necessity, must also make such decisions, while Pareto-agnostic concepts, by design, do not.

I conclude this section by considering a policymaker who is reluctant to hurt a particular agent in order to help others—a situation that might arise if the agent to be hurt were a member of a socially disadvantaged group. That policymakers often report separate rank distributions for different classes of agents suggests this situation is empirically relevant. (Recall from the introduction that the SFUSD reports the rank distribution by ethnicity and zipcode, and the NRMP reports match rates by whether the graduate's medical school is allopathic or osteopathic and by whether it is foreign or domestic.) To codify such policymaker preferences, I introduce a family of efficiency concepts that only consider tradeoffs within defined sets of agents. This family bridges the gap between ordinal efficiency and rank efficiency in a natural way.

3.1 Ex Post and Ordinal Efficiency

I BEGIN by briefly reviewing ex post and ordinal efficiency. Hold fixed an ordinal market. A deterministic assignment, x , is said to **ex post dominate** another deterministic assignment, y , if all agents rank their allocations in x at least as high as their allocations in y , with at least one agent's preference strict. If a deterministic assignment is not ex post dominated by any other deterministic assignment, it is **ex post efficient**.

In an ordinal setting, there is an ambiguity in extending this definition to nondeterministic assignments. How can the planner tell whether an agent prefers one lottery over object types to another without knowing her cardinal utility? There are two ways forward. The first is to require that the

planner is never forced into implementing an ex post-inefficient, deterministic assignment. Formally, the definition above can be extended in this way by calling a (potentially nondeterministic) assignment **ex post efficient** if it has an implementation whose realizations are all ex post efficient.⁹

The second way forward is to assert an agent’s preference between two lotteries if it would hold regardless of the cardinal utilities that rationalize her ordinal ranking. An assignment, x , is said to **ordinally dominate** another assignment, y , if all agents weakly prefer (in the sense of first-order stochastic dominance) their allocations in x to their allocations in y , with at least one agent’s preference strict. To formalize this concept, define the **personal rank distribution of agent a in assignment x** by $N^x(k;a) \equiv \sum_o \mathbb{1}\{r_{ao} \leq k\} x_{ao}$. Intuitively, $N^x(k;a)$ represents the probability in assignment x that agent a gets an object type she ranks k th or better. Mathematically, assignment x ordinally dominates assignment y if $N^x(k;a) \geq N^y(k;a)$ for all agents, a , and all ranks, k , with the inequality strict for at least one agent–rank pair. If an assignment is not ordinally dominated by any other, it is **ordinally efficient**.

The concepts described thus far can be illustrated with the simple four-agent–four-object-type example in [Figure 2](#), where each object type has unit capacity. The profile of rankings is listed in part (a). While all agents prefer objects o_1 and o_2 to the other objects, agents a_1 and a_2 rank object o_1 first, while agents a_3 and a_4 rank object o_2 first. Intuitively, a policymaker might want to give objects o_1 and o_2 to agents that rank them first; however, this is not required by ex post efficiency—see the deterministic, ex post-efficient assignments w and x in part (b).¹⁰

This defect is especially troubling when considering nondeterministic, ex post-efficient assignments, like the 50/50 lottery over assignments w and x that is listed in part (c) as $\frac{1}{2}(w + x)$. Thinking in terms of probabilities, agent a_2 would gladly trade her half-share of object o_2 (agent a_2 ’s 2nd choice) for agent a_3 ’s half-share of object o_1 (agent a_2 ’s 1st choice). Agent a_3 is also strictly better off after this trade. So, how can assignment $\frac{1}{2}(w + x)$ be ex post efficient? The answer is simple: the trade in probability shares just described is impossible ex post. (When the realized deterministic assignment is w , agent a_2 gets object o_2 , but agent a_3 doesn’t get object o_1 , while when the realized deterministic assignment is x , agent a_3 gets object o_1 , but agent a_2 doesn’t get object o_2 .)

Ordinal efficiency solves this problem by considering interim trades—

-
9. Although unintuitive, it is possible for an ex post efficient assignment to have two different implementations, one of whose support is entirely ex post efficient and one of whose is not. (See Example 2 in Abdulkadiroğlu and Sönmez 2003.) Rank efficient and ordinally efficient assignments, which will be defined later in this section, do not have this problem.
 10. By [Proposition 4](#) in the next section, any serial dictatorship yields an ex post efficient, deterministic assignment. Assignment w is ex post efficient because it is the outcome of the serial dictatorship where agents choose in the order a_1 — a_2 — a_3 — a_4 . The order that corresponds to assignment x is a_3 — a_4 — a_2 — a_1 .

$$\begin{aligned} \mathcal{A} &= \{a_1, a_2, a_3, a_4\} \\ \mathcal{O} &= \{o_1, o_2, o_3, o_4\} \\ q &= \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \end{aligned} \quad r = \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} 1 & 2 & 3 & 4 \end{pmatrix} \\ a_2 & \begin{pmatrix} 1 & 2 & 3 & 4 \end{pmatrix} \\ a_3 & \begin{pmatrix} 2 & 1 & 3 & 4 \end{pmatrix} \\ a_4 & \begin{pmatrix} 2 & 1 & 4 & 3 \end{pmatrix} \end{pmatrix}$$

(a) Agents (\mathcal{A}), Object Types (\mathcal{O}), Capacities (q), and Rankings (r)

$$\begin{aligned} w &\equiv \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix} \\ a_2 & \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix} \end{pmatrix} & x &\equiv \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix} \\ a_2 & \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix} \end{pmatrix} & \frac{1}{2}(w+x) &= \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{pmatrix} \\ a_2 & \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix} \end{pmatrix} \\ y &\equiv \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix} \\ a_2 & \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix} \end{pmatrix} & z &\equiv \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix} \\ a_2 & \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} 0 & 0 & 1 & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix} \end{pmatrix} & \frac{1}{2}(y+z) &= \begin{pmatrix} o_1 & o_2 & o_3 & o_4 \\ a_1 & \begin{pmatrix} \frac{1}{2} & 0 & 0 & \frac{1}{2} \end{pmatrix} \\ a_2 & \begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \end{pmatrix} \\ a_3 & \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix} \\ a_4 & \begin{pmatrix} 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix} \end{pmatrix} \end{aligned}$$

(b) Some Ex Post-Efficient,
Deterministic Assignments

(c) Some Ex Post-Efficient,
Nondeterministic Assignments

EXPECTED NUMBER OF AGENTS WHO GET...				
	1st choice	2nd or better	3rd or better	4th or better
w	1	2	4	4
x	1	2	3	4
$\frac{1}{2}(w+x)$	1	2	$3\frac{1}{2}$	4
y	2	2	4	4
z	2	2	3	4
$\frac{1}{2}(y+z)$	2	2	$3\frac{1}{2}$	4

(d) Rank Distributions for Assignments in Parts (b) and (c)

Figure 2: Example Illustrating Ex Post, Ordinal, and Rank Efficiency Concepts

NOTES: All assignments in this figure except $\frac{1}{2}(w+x)$ are ordinally efficient; however, of the assignments in this figure, only y is rank efficient, since it attains the best possible rank distribution in this ordinal market. (Everyone gets a 3rd-or-better allocation, and it is impossible to give more than two 1st-choice or 2nd-or-better allocations, since only two objects are ranked 1st or 2nd.)

exchanges in probability shares that occur before a specific deterministic assignment is realized. Consider the 50/50 lottery over assignments y and z that is listed in part (c) as $\frac{1}{2}(y+z)$. Clearly, it ordinally dominates $\frac{1}{2}(w+x)$, since agents a_2 and a_3 achieve the trade discussed in the previous paragraph.¹¹ But how was this trade implemented? Relative to assignment w , in assignment y , agent a_2 is giving up her 2nd choice for her 3rd choice, while agent a_3 is getting his 1st choice in exchange for his 3rd choice. The situation is reversed when considering assignment z relative to assignment x . Essentially, agent a_2 is agreeing to take a hit when the realization is y in exchange for agent a_3 agreeing to take a hit when the realization is z .

By allowing a wider variety of trades, ordinal efficiency is necessarily a refinement of ex post efficiency on the domain of all assignments. But, on the domain of deterministic assignments, ordinal efficiency fails to have any bite relative to ex post efficiency, since deterministic assignments have no interim stage. Formally,

Proposition 1 (Bogomolnaia and Moulin 2001, Parts a and b).

- (a) *For all ordinal markets, an ordinally efficient assignment must also be ex post efficient.*
- (b) *For some ordinal markets, there are ex post-efficient assignments that fail to be ordinally efficient.*
- (c) *For all ordinal markets, a deterministic assignment is ex post efficient if and only if it is ordinally efficient.*

As noted, parts (a) and (b) have been previously shown by Bogomolnaia and Moulin (2001). Part (c), although straightforward, does not appear in previous literature. Its implication for practice is worth mentioning: for a policymaker who cares only about the realized deterministic assignment, ordinal efficiency doesn't offer anything new over ex post efficiency, while rank efficiency does (as I will show in Proposition 2). Note that I have combined parts (a)–(c) into one proposition so that it retains a parallel structure with the upcoming Proposition 2.

So, although there is a strong economic intuition for thinking in terms of ordinal efficiency instead of ex post efficiency, to a policymaker who cares only about the realized deterministic assignment, there is no difference between the two concepts. In the next section, I present a refinement of ordinal efficiency that even the policymaker just mentioned would recognize.

3.2 Rank Efficiency

ORDINAL EFFICIENCY captured an important improvement for the example illustrated in Figure 2, but it left another unaddressed. While agents a_1 , a_2 , and a_3 all rank object o_4 last, agent a_4 only ranks it second to last. Since

11. In fact, assignment $\frac{1}{2}(y+z)$ is ordinally efficient. This follows readily from Theorem 1 of Bogomolnaia and Moulin (2001).

some agent must end up with object o_4 , it is intuitive for a policymaker to prefer that agent to be a_4 .

Still holding the ordinal market fixed, define the **rank distribution of assignment x** , denoted by $N^x(k)$, to be the sum of all agents' personal rank distributions in x , that is, $N^x(k) \equiv \sum_a N^x(k; a)$. Intuitively, $N^x(k)$ is the expected number of agents who receive their k th choice or better in assignment x —a fact that is made more obvious by plugging in the definition of the personal rank distribution to yield

$$N^x(k) = \sum_a \sum_o \mathbb{1}\{r_{ao} \leq k\} x_{ao}.$$

An assignment, x , is said to **rank-dominate** another assignment, y , if the rank distribution of x first-order stochastically dominates that of y , that is, if $N^x(k) \geq N^y(k)$ for all ranks, k , with the inequality strict for at least one k . If an assignment is not rank dominated by any other, it is **rank efficient**.

For the example in [Figure 2](#), rank distributions for the assignments in parts (b) and (c) are tabulated in part (d). Assignment y rank-dominates all other assignments in the figure, which lines up with the intuition that agent a_4 should get object o_4 , since she dislikes it less than everyone else. It also shows that for a policymaker who cares only about the realized deterministic allocation, rank efficiency is distinct from ex post efficiency, since the ex post-efficient, deterministic assignments w , x , and z in part (b) are all rank dominated by the deterministic assignment y .

Although rank efficiency may look different than the more standard efficiency concepts of [Section 3.1](#), it is actually a refinement of them—a result that follows immediately from the rank distribution being the sum of all agents' personal rank distributions. Improving all agents' personal rank distributions must also improve their sum. Formally,

Proposition 2.

- (a) *For all ordinal markets, a rank-efficient assignment must also be ordinally efficient.*
- (b) *For some ordinal markets, there are ordinally efficient assignments that fail to be rank efficient.*
- (c) *Statement (b) continues to hold if the assignment must be deterministic.*

Parts (a) and (b) show that rank efficiency rejects the ordinally efficient assignments whose rank distributions can be stochastically dominated. Comparing to part (c) of [Proposition 1](#), part (c) of [Proposition 2](#) suggests that, to a policymaker who cares only about deterministic assignments, rank efficiency is distinct from ex post efficiency while ordinal efficiency is not.

3.3 Bridging the Gap Between Ordinal and Rank Efficiency

IN THE EXAMPLE from [Figure 2](#), consider moving from assignment w to assignment y . Agent a_2 moves from her 2nd choice to her 3rd, while agent

a_3 moves from his 3rd choice to his 1st. Ordinal efficiency remains agnostic about this trade, since it hurts the former agent to help the latter. On the other hand, rank efficiency insists on the trade, in some sense treating the two agents' utilities as comparable (a sense that will be greatly clarified in [Section 4.2](#)). But what if the policymaker were uncomfortable with such a comparison, perhaps because the agent being hurt is a member of a socially disadvantaged group? That real-world policymakers often report separate rank distributions for different classes of agents suggests this concern is empirically relevant. (See the beginning of [Section 3](#) for examples.)

I capture which agent utility comparisons are valid with a **partition** of the agents, that is, a collection of non-empty, disjoint sets whose union is \mathcal{A} . For instance, the San Francisco Unified School District (SFUSD) reports rank distributions by ethnicity; the corresponding partition would have one element for each reported ethnicity. To capture the comparability of agent utilities within a partition element, I define the **rank distribution of the set of agents \mathcal{A}_i** to be the sum of the personal rank distributions of all agents in \mathcal{A}_i , that is, $N^x(k; \mathcal{A}_i) \equiv \sum_{a \in \mathcal{A}_i} N^x(k; a)$.

Using these definitions, I call one assignment superior to another if its rank distributions are weakly stochastically better on all elements of the partition, with strict stochastic dominance holding for at least one. Formally, given a partition, $\mathcal{P} = \{\mathcal{A}_i\}$, an assignment, y , is **rank-dominated with respect to \mathcal{P}** (or \mathcal{P} -rank dominated, for short) by another assignment, x , if $N^x(k; \mathcal{A}_i) \geq N^y(k; \mathcal{A}_i)$ for all partition members, \mathcal{A}_i , and ranks, k , with the inequality strict for at least one partition member–rank pair. If there is no other assignment that rank-dominates x with respect to a partition, \mathcal{P} , then x is **partitioned rank efficient with respect to \mathcal{P}** (or \mathcal{P} -rank efficient, for short). In the context of this partition-based concept, ordinal efficiency and rank efficiency lie at the extremes: the former corresponds to the finest partition of agents while the latter corresponds to the coarsest.

For the sake of concreteness, consider the example in [Figure 2](#) with two partitions: $\mathcal{P}_1 = \{\{a_1, a_4\}, \{a_2, a_3\}\}$ and $\mathcal{P}_2 = \{\{a_1, a_2\}, \{a_3, a_4\}\}$. In part (b), assignment w is \mathcal{P}_1 -rank dominated by assignment y , but fails to be \mathcal{P}_2 -rank dominated by it. The intuition is simple: the “improvement” in going from assignment w to assignment y involves hurting agent a_2 to help agent a_3 . This improves both of partition \mathcal{P}_1 's rank distributions but worsens one of \mathcal{P}_2 's rank distributions.

The family of partitioned rank efficiency concepts allows a generalization of [Proposition 2](#) involving **partition refinements**, where partition \mathcal{P}' is said to refine partition \mathcal{P} if the two partitions are distinct and every element of \mathcal{P}' is contained in an element of \mathcal{P} . That is, \mathcal{P}' refines \mathcal{P} if $\mathcal{P}' \neq \mathcal{P}$ and for each \mathcal{A}'_i in \mathcal{P}' , there is a \mathcal{A}_i in \mathcal{P} such that $\mathcal{A}'_i \subseteq \mathcal{A}_i$.

Proposition 3.

- (a) *For all ordinal markets, for all partitions of the agent set, \mathcal{P} and \mathcal{P}' , where the latter refines the former, a \mathcal{P} -rank-efficient assignment must also be \mathcal{P}' -rank ef-*

icient.

- (b) Fix an agent set, \mathcal{A} , with three or more agents, and let \mathcal{P} and \mathcal{P}' be partitions of that set where the latter refines the former.

Then, for some ordinal markets with agent set \mathcal{A} , there are \mathcal{P}' -rank-efficient assignments that fail to be \mathcal{P} -rank efficient.

- (c) Statement (b) continues to hold if the assignment must be deterministic.

Note that statement (b)'s requirement on the number of agents only serves to rule out a few uninteresting scenarios.¹²

In the next section, I will show how the specific partition used to define a partitioned rank efficiency concept corresponds to assumptions about the (unknown to the planner) cardinal utilities underlying agents' rankings.

4 Mechanisms

IN THIS section I will begin by reviewing mechanisms connected to the efficiency concepts from previous literature discussed in [Section 3](#). Allowing agents to choose objects in some specified order (i.e., serial dictatorship), is tightly connected to ex post efficiency, and indeed, a randomized form of serial dictatorship is ubiquitous in the field. Similarly, the “simultaneous eating” mechanism of Bogomolnaia and Moulin (2001) is tightly connected to ordinal efficiency; however, it is absent in the field, which is puzzling given the intuitive appeal of ordinal efficiency as an economic concept.

One potential resolution to the puzzle lies in the asymptotic equivalence result of Che and Kojima (2010), which shows that, in certain large markets, random serial dictatorship yields an almost identical outcome to that of the most natural simultaneous eating mechanism, the “probabilistic serial” mechanism. However, the equivalence ceases to hold in asymptotics where the number of agents fails to grow large relative to the number of object types (Manea 2009).

An alternative resolution to the puzzling absence of simultaneous eating mechanisms in the field comes from noticing that there is another, more widely used, class of linear programming (LP) mechanisms that are just as tightly tied to ordinal efficiency. Each such mechanism takes the submitted rankings, makes assumptions about the underlying cardinal utilities, and then maximizes a welfare sum.¹³ Similar mechanisms have been seen in

-
12. With only one agent, statement (b) is vacuous, since there is only one possible partition and the definition of refinement requires two. When there are two agents—call them a_1 and a_2 —the only possibility for the partitions is $\mathcal{P} = \{\{a_1, a_2\}\}$ and $\mathcal{P}' = \{\{a_1\}, \{a_2\}\}$. If it is possible to give both a_1 and a_2 their respective first choices, then this is the only \mathcal{P}' -rank-efficient assignment, and it can't possibly be \mathcal{P} -rank dominated. Otherwise, the only \mathcal{P}' -rank-efficient assignment involves giving one agent a first choice and one a second choice, and again it can't possibly be \mathcal{P} -rank dominated.
13. Maximizing a welfare sum can be interpreted as maximizing the expected utility of an agent who is behind the Rawlsian “veil of ignorance” (Harsanyi 1955, 1975, 1986; Rawls 1972).

the field, such as the match of house officers to hospitals in Cambridge and London and the TFA match mentioned in the introduction. The mechanism can also be implemented via a “tinkering” process that I will introduce in [Section 5](#).

I will formally define this class of welfare-maximizing mechanisms and then show that a common-sense restriction on the assumptions made about the cardinal preferences leads to a family of mechanisms that is intimately connected to rank efficiency. Each such mechanism is quite intuitive. First, it scores each possible assignment, giving it some number of points for each agent who gets her top choice, some smaller number of points for each agent who gets her second choice, and so on. It then selects an assignment that maximizes that score.

I conclude this section by relating assumptions on the underlying cardinal utilities to the family of partitioned rank efficiency concepts from the previous section. Doing so provides a unifying framework for many efficiency concepts used in the matching literature.

4.1 Random Serial Dictatorship

DUE TO their familiarity, I begin by clarifying the “tight” connection between randomized serial dictatorship and ex post efficiency that I mentioned above. When I present similar results for the other concepts introduced in [Section 3](#), the parallel will be clear.

Serial dictatorships simply let agents choose in some specified order. Formally, define a **deterministic ordering of agents** to be a one-to-one mapping, $\pi : \{1, \dots, |\mathcal{A}|\} \mapsto \mathcal{A}$. Then, **serial dictatorship with respect to deterministic ordering π** , denoted SD_π , is the ordinal assignment mechanism that selects the deterministic assignment where agent $\pi(1)$ gets a copy of the $r_{\pi(1)}$ -maximal object type, agent $\pi(2)$ gets a copy of the $r_{\pi(2)}$ -maximal object type among those that still have copies remaining, and so on.

Fairness sometimes dictates that the ordering be randomized. Define a **random ordering of agents** to be a random variable whose realized values are deterministic orderings of agents. Then, **random serial dictatorship with respect to random ordering Π** , denoted RSD_Π , is the ordinal assignment mechanism that selects the expectation of the deterministic assignment chosen by serial dictatorship with respect to the realized deterministic ordering of agents. That is, for any ordinal market, \mathcal{M} , I define $RSD_\Pi(\mathcal{M}) \equiv \mathbb{E}[SD_\Pi(\mathcal{M})]$, where all randomness inside the expectation comes from Π . Now, the tight connection between ex post efficiency and random serial dictatorship can be made precise.

Proposition 4 (Svensson 1994). *Fix an ordinal market, \mathcal{M} . An assignment, x , is ex post efficient if and only if there exists some random ordering of agents, Π , such that $x = RSD_\Pi(\mathcal{M})$.*

As noted, [Proposition 4](#) was proven in the deterministic setting by Svensson (1994). (See also Abdulkadiroğlu and Sönmez 1998 and Manea 2007.)

In the sections that follow, I present analogues of [Proposition 4](#) for ordinal efficiency, rank efficiency, and the partitioned rank efficiency concepts that lie between.

4.2 The Welfare Maximization Mechanism

Now, I introduce a class of mechanisms that make assumptions about the cardinal utilities underlying ordinal rankings. To model such an assumption, I will use a **valuation vector**—an $|\mathcal{O}|$ -dimensional vector that is **strictly decreasing** (i.e., its first component is strictly greater than its second, and so on) and whose $|\mathcal{O}|$ th component is zero. A valuation vector’s k th component can be interpreted as the von Neumann–Morgenstern utility associated with getting a k th-choice object type. Agent valuations can be stacked into an $|\mathcal{A}| \times |\mathcal{O}|$ **valuation matrix**, ω , whose (a, k) entry represents the cardinal utility that agent a places on getting her k th choice. Given the set of assumptions codified by a valuation matrix, it is natural to choose an assignment that maximizes the welfare sum. Formally, the **welfare maximization mechanism with respect to valuation matrix ω** (or ω -welfare maximization mechanism for short), denoted WM_ω , maps any ordinal market, $(\mathcal{A}, \mathcal{O}, r, q)$, to the set of assignments given by

$$\begin{aligned} \arg \max_{(x_{ao}) \geq 0} \quad & \sum_a \sum_o \omega_{ar_{ao}} x_{ao}, \\ \text{s.t.} \quad & \sum_a x_{ao} \leq q_o, \quad \text{for all } o \in \mathcal{O}, \\ & \sum_o x_{ao} = 1, \quad \text{for all } a \in \mathcal{A}. \end{aligned} \tag{1}$$

The objective of this linear program is the **welfare objective with respect to valuation matrix ω** (or the ω -welfare objective, for short). Technically, for WM_ω to be an ordinal assignment mechanism (as defined in [Section 2](#)), I need a selection rule to deal with cases where the arg max isn’t a singleton; however, the specifics of this rule won’t matter until I consider incentives. As such, I defer discussion of this detail to [Section 6.2.1](#).

With no further restrictions on the valuation matrix, the welfare maximization mechanism has the same tight relationship with ordinal efficiency that random serial dictatorship has with ex post efficiency. Formally,

Proposition 5 (McLennan 2002; Manea 2008). *Fix an ordinal market, \mathcal{M} . An assignment, x , is ordinally efficient if and only if there exists some valuation matrix, ω , such that $x \in WM_\omega(\mathcal{M})$.*

As noted, this result has previously been shown by McLennan (2002) and Manea (2008). It follows as a special case of the more general [Proposition 6](#), which I introduce below.

So far, I have not placed any strong restrictions on the valuation matrix. This affords considerable freedom to the policymaker—in principle, she could assume a different valuation vector for each agent. With detailed

information about the agents, this freedom might be useful, but such knowledge is often not available. When two agents look the same to the policymaker, it is sensible to constrain her to assume that they both have the same valuation vector. As I will show, such constraints correspond to the efficiency concepts introduced in [Section 3](#), which in turn correspond to policymaker behavior that is observed in the field.

To capture which agents “look the same”, I again use a partition of the agents. A valuation matrix, ω , **respects the partition \mathcal{P}** (or is \mathcal{P} -respecting, for short) if any two agents in the same partition element have the same valuation, that is, if two agents, a and a' , are both in some element, \mathcal{A}_i , of \mathcal{P} , then $\omega_{ak} = \omega_{a'k}$ for all ranks, k . For a \mathcal{P} -respecting valuation matrix, ω , I denote the valuation vector for agents in a partition element, \mathcal{A}_i , by $\omega_{\mathcal{A}_i}$.

Given a partition, \mathcal{P} , the connection between welfare maximization mechanisms that use a \mathcal{P} -respecting valuation matrix and \mathcal{P} -rank efficiency is strong. Formally,

Proposition 6. *Fix an ordinal market, \mathcal{M} , and a partition of the agents in that market, \mathcal{P} . An assignment, x , is \mathcal{P} -rank efficient if and only if there exists some \mathcal{P} -respecting valuation matrix, ω , such that $x \in \text{WM}_\omega(\mathcal{M})$.*

The proof places assignments in the Euclidean space where each axis represents the number of agents in a particular partition element who got their k th choice or better. In this space, the ω -welfare objective can be more intuitively written in terms of the partitioned rank distributions as

$$\sum_{i=1}^{|\mathcal{P}|} \sum_{k=1}^{|\mathcal{O}|-1} [\omega_{\mathcal{A}_i k} - \omega_{\mathcal{A}_i (k+1)}] N^x(k; \mathcal{A}_i). \quad (2)$$

Since the square-bracketed terms are strictly positive, necessity is clear. Sufficiency comes from a supporting hyperplane argument, although some convex analysis subtleties must be addressed to ensure that the supporting hyperplane corresponds to a vector that is strictly decreasing, i.e., a valuation.

4.3 The Rank-Value Mechanism

THE EXTREME elements of the set of partitioned rank efficiency concepts are the most interesting. As I just showed, ordinal efficiency—or partitioned rank efficiency with respect to the finest partition of agents—is the result of a welfare maximization process where assumptions about the cardinal utilities that rationalize ordinal rankings are unconstrained. On the other end of the spectrum, rank efficiency—or partitioned rank efficiency with respect to the coarsest partition of agents—constrains the policymaker to assume that all agents can be modeled with the same valuation vector.

Going forward, I will focus on rank efficiency, since the global rank distribution is a more common concern than the partitioned one. Towards this end, the **rank-value mechanism with respect to valuation vector v** (or

v -rank-value mechanism, for short), denoted RV_v , maps any ordinal market to the set returned by the welfare maximization mechanism with respect to a valuation matrix where every row is v . That is, if $\omega_{\alpha k} = v_k$ for any agent, α , and rank, k , then $RV_v \equiv WM_\omega$.¹⁴ When ω obeys the conditions in the previous sentence, the ω -welfare objective can also be referred to as the **welfare objective with respect to v** (or the v -welfare objective, for short). To be clear, the assumptions on ω do not require that all agents have the same ranking, but rather that they share a common cardinal preference profile in going from first choice to second choice, and so on. The following is a clear corollary to [Proposition 6](#).

Proposition 7. *Fix an ordinal market, \mathcal{M} . An assignment, x , is rank efficient if and only if there exists some valuation vector, v , such that $x \in RV_v(\mathcal{M})$.*

Before moving on, I emphasize two important points concerning the interpretation of a rank-value mechanism’s valuation vector. First, it need not be seen as the common cardinal utility profile for all agents. If agent utility profiles are realizations of identically distributed random vectors, then one can think of the k th component of the valuation vector as the *expected* cardinal utility of a k th choice. In other words, the rank-value mechanism can be thought of as maximizing *expected* welfare when the expectation of the agents’ true cardinal utility profiles is known, but their realizations are not.

Second, although I have framed the discussion in this section around welfare sums, the policymaker need not think about the valuation vector in this way. For instance, in a job-allocation setting, v_k could codify the policymaker’s assumption about the probability that an agent will show up to a job that she ranked k th. In this particular example, the rank-value mechanism would find an assignment that maximizes the expected number of agents who show up for their assigned jobs. (See [Davis and Montagnes 2020](#), which considers a version of this approach in which the match-value function is perfectly known by the agents, but imperfectly known by the planner.)

14. One might think the Boston mechanism ([Abdulkadiroğlu and Sönmez 2003](#); [Kojima and Ünver 2014](#)) without priorities is equivalent to a rank-value mechanism where the valuation vector is **near lexicographic**, that is, where $v_1 \gg v_2 \gg \dots \gg v_{|\mathcal{O}|}$. This is not the case, since the Boston mechanism’s greedy approach can easily fail to find the solution to the rank-value mechanism’s defining LP. Intuitively, consider an ordinal market in which there are two assignments with the maximum number of first-choice allocations, and further, let one of these assignments give more second-choice allocations than the other. Boston will choose between the two assignments at random, while a near-lexicographic rank-value mechanism will always choose the assignment with more second-choice allocations.

However, looking to the computer-science literature, a rank-value mechanism with a near-lexicographic valuation vector does yield the set of rank-maximal matchings ([Irving et al. 2006](#); [Kavitha and Shah 2006](#); [Paluch 2013](#)).

Efficiency Concept	ex post	\Leftarrow \nRightarrow	ordinal	\Leftarrow \nRightarrow	partitioned rank	\Leftarrow \nRightarrow	rank
Level of Rank Distribution Aggregation	—		agent		partition element		market
Restriction on Valuation Assumptions	—		none		agents in a partition element must have the same valuation		all agents must have the same valuation

Table 1: Framework Relating Efficiency Concepts

NOTES: In the top row, arrows indicate logical implications. In the middle row, for each efficiency concept, the level of rank distribution used to define it is listed. In the bottom row, for each efficiency concept, the restriction on the valuation matrix required to make the welfare maximization mechanism satisfy that concept is listed. There are two empty entries under ex post efficiency because that concept is not defined in terms of a rank distribution.

4.4 A Unifying Framework

When combined, the results from Sections 3 and 4 provide a unifying framework for many of the efficiency concepts in the matching literature. This framework is sketched in Table 1. The common thread: when an efficiency concept aggregates a group’s rank distribution, the corresponding welfare maximization mechanism must use the same valuation for all members of that group. This simple idea spans an entire family of efficiency concepts—running from ordinal efficiency to rank efficiency—that can be implemented by solving a simple linear program that makes an assumption about the cardinal utilities and then finds a welfare-maximizing assignment. The presence of such LP mechanisms in the field (see, e.g., Roth 1991 and Ünver 2001, 2005) shows that this family of efficiency concepts is empirically relevant.

5 A Model of Tinkering

AN ASSIGNMENT is rank efficient if and only if it is generated by the rank-value mechanism for some valuation vector. While this is theoretically interesting, one might wonder whether rank efficiency actually describes a real-world phenomenon. After all, computing the solution to a linear program involving a welfare sum might seem beyond the technical sophistication of many policymakers. In this section, I argue against that claim by showing that rank-efficient assignments are the natural outcome of a simple sequential improvement process that starts with an arbitrary deterministic

assignment and looks for changes that locally improve the welfare objective. In this sense, no knowledge of linear programming is required to solve the welfare maximization LP laid out in [Equation 1](#).

This process models a potentially widespread phenomenon, which I call **tinkering**. Although clearinghouses are usually set up by market designers, they often end up being administered by people who might not understand the importance of running a mechanism exactly as it was designed. If these administrators see a way to “improve” on the assignment outputted by a mechanism—to *tinker* with it—they seem likely to do so unless specifically convinced otherwise.¹⁵ As described in the introduction, this is precisely how I observed TFA matching teachers to regions in 2010. (See footnote 2.)

To clarify the idea of tinkering, consider an administrator in the world of [Figure 2](#) who is in charge of executing a computer program that randomly picks an ordering of the four agents from the uniform distribution and then runs serial dictatorship. By [Proposition 4](#), any deterministic, ex post-efficient assignment could be selected. Imagine the program selects assignment w . A market designer knows that to get strategyproofness, sometimes assignments like w must be allowed. But the administrator might want to switch the allocations of agents a_2 and a_3 —yielding assignment y —since doing so would stochastically improve the rank distribution. In fact, one could even imagine the administrator blithely re-running the mechanism (with a new ordering of agents each time) until it returns assignment y , not realizing that doing so undermines carefully designed incentive properties.

Of course, tinkering might not be widespread if there isn’t much leeway for the rank distribution to be improved. I show that this was not the case for the Harvard Business School (HBS) Field Immersion Experiences for Leadership Development (FIELD) match that I helped to design. (See footnote 1.) Using rankings submitted by the 900 MBAs who participated in that match, I compute the expected rank distribution under random serial dictatorship (the mechanism used by HBS) and then compare it to the counterfactual rank distributions produced by rank-value mechanisms for several valuations. The rank-value mechanisms increase the number of students who get a second-or-better ranked country by roughly 18 percent. In light of this gain, it is not surprising that HBS briefly considered ex post tinkering until they were reminded that doing so would undermine the strategyproofness that was advertised to students. I discuss issues surrounding truthful preference revelation in [Section 6](#).

If the potential gains from tinkering are similarly large in other markets, one should expect that the temptation to tinker is strong, especially behind closed doors. A market designer could implement a strategyproof mech-

15. Such tinkering is explicitly carried out by a computer algorithm in the recently adopted procedure used by the Israeli Ministry of Health to match newly minted physicians to internships in that country. See Bronfman et al. (2015).

anism, only to have the policymaker tinker it into a rank-efficient mechanism.¹⁶ As such, a formal model of tinkering is empirically relevant. In the next subsection, I introduce a model in which the policymaker tinkers until there are no welfare improvements left to be made.¹⁷

5.1 Sequential Improvements

TO FORMALIZE the tinkering process described above, I require notation for allocation swaps. Fix an ordinal market. A **trade cycle in assignment x** is a list of agent–object type pairs, $((a_i, o_i))_{i=1}^n$, where each agent is in at most one pair and holds a positive amount of her paired object type (that is, $x_{a_i o_i} > 0$ for all i in $\{1, \dots, n\}$). It represents the idea that agent a_1 passes a share of o_1 to agent a_2 , who in turn passes a share of o_2 to agent a_3 , and so on, until agent a_n passes a share of object type o_n to agent a_1 , closing the cycle. The **capacity of a trade cycle** is the maximum share that can be passed in the way just described, that is, $\min_{i \in \{1, \dots, n\}} x_{a_i o_i}$. **Executing a trade cycle in assignment x yields assignment y** if starting with x and implementing the trades in the cycle as described above, up to capacity, yields y . That is, if the trade cycle’s capacity is c , then assignment y is the same as x , except $y_{a_i o_i} = x_{a_i o_i} - c$ and $y_{a_i o_{i-1}} = x_{a_i o_{i-1}} + c$, for all i in $\{1, \dots, n\}$, where o_0 is defined to be o_n .

Of course, a trade cycle can’t introduce unallocated shares into the system. A **claim chain in assignment x** is an object type and a list of agent–object type pairs, $(o_s, ((a_i, o_i))_{i=1}^n)$, where each agent is in at most one pair and holds a positive amount of her paired object type (that is, $x_{a_i o_i} > 0$ for all i in $\{1, \dots, n\}$), and where there is some unallocated amount of object type o_s (that is, $\sum_a x_{a o_s} < q_{o_s}$). It represents the idea that agent a_1 claims an unallocated share of object type o_s and then passes a share of o_1 to agent a_2 , who in turn passes a share of o_2 to agent a_3 , and so on, until agent a_n simply deallocates a share of object type o_n , ending the chain. A claim chain’s capacity and execution are defined as one would expect. The **capacity of a claim chain** is the minimum of $\min_{i \in \{1, \dots, n\}} x_{a_i o_i}$ and $q_{o_s} - \sum_a x_{a o_s}$, where the second expression represents the amount of o_s that is unallocated in assignment x . If the capacity of a claim chain is c , then **executing that claim chain in assignment x yields y** if assignment y is the same as assignment

-
16. One might worry that this tinkering temptation would be mitigated over time as agents learn to deviate from truth-telling (as [Proposition 10](#) suggests they might). Pycia (2019) shows that for a given true preference profile, anonymous statistics (such as the rank distribution) are asymptotically equivalent across all strategyproof, Pareto-efficient mechanisms. This result extends to complete-information Nash equilibria of non-strategyproof mechanisms; however, it does not readily extend to incomplete-information environments. Lab subjects also generally fail to converge to non-truth-telling equilibria in a matching setting (see, e.g., Featherstone and Niederle 2016).
17. One might also consider a richer behavioral model in which not all potential welfare improvements are made due to some sort of tinkering cost. The simpler model presented in this section corresponds to the limit in which this cost is zero.

x , except $y_{a_i o_i} = x_{a_i o_i} - c$ and $y_{a_i o_{i-1}} = x_{a_i o_{i-1}} + c$, for all i in $\{1, \dots, n\}$, where o_0 is defined to be o_s .

Finally, I introduce the method for determining whether the good from executing a trade cycle or claim chain outweighs the bad. Given a valuation vector, v , define the **v -welfare change** for a trade cycle $\tau = ((a_i, o_i))_{i=1}^n$ to be

$$\Delta(\tau; v) \equiv \sum_{i=1}^n \left(v_{r_{a_i o_{i-1}}} - v_{r_{a_i o_i}} \right), \quad (3)$$

where o_0 is defined to be o_n . The v -welfare change for a claim chain, $(o_s, ((a_i, o_i))_{i=1}^n)$, can be similarly defined by letting o_0 represent o_s instead of o_n in Equation 3. Intuitively, each term in the expression represents the (per-unit-traded) utility gain of an agent, given the assumptions encoded by the valuation vector. A trade cycle or claim chain, σ , is **v -welfare improving** if $\Delta(\sigma; v)$ is strictly positive.

It turns out that v -welfare-improving trade cycles and claim chains are tightly connected to rank efficiency. Formally,

Proposition 8. *Fix an ordinal market. An assignment, x , is rank efficient if and only if there exists a valuation vector, v , such that there is no trade cycle or claim chain in x that is v -welfare improving.*

Given that executing a v -welfare-improving trade cycle or claim chain would increase the v -welfare objective, necessity is straightforward. Sufficiency hinges on an interesting fact (proven as Lemma A.3 in Appendix A): given any initial and final assignments, there exists a set of trade cycles and claim chains in the initial assignment that can be executed in any order to reach the final assignment.¹⁸ If the v -welfare objective of the final assignment is greater than that of the initial assignment, then at least one of the transitional trade cycles or claim chains must be v -welfare improving.

One might think that the condition in Proposition 8 should involve trade cycles and claim chains that *stochastically improve* the rank distribution when executed. Such a trade cycle or claim chain, which I will call **rank improving**, is one that is v -welfare improving for all valuations, v . It turns out that the absence of rank-improving trade cycles and claim chains is necessary, but not sufficient, for rank efficiency. (See Proposition A.1 in Appendix A.) This stands in contrast to the characterization of ordinal efficiency as the absence of any Pareto-improving trades in probability shares. (See Lemma 3 in Bogomolnaia and Moulin 2001.) Intuitively, a rank improvement might decompose into a set of trade cycles and claim chains that only stochastically improve the rank distribution when considered together.¹⁹

18. Here, I sweep an important detail under the rug: the “volume” pushed along a given trade cycle or claim chain (c in the definitions above) might need to be less than full capacity. For a formal discussion, see Section A.3 in Appendix A.

19. The proof to Proposition A.1 in Appendix A presents an example (listed in Figure A.1) of a

INPUT: a deterministic assignment, x^0 , and a valuation vector, v ;

$t \leftarrow 0$;

while x^t has a v -welfare-improving trade cycle or claim chain

σ^t	\leftarrow any v -welfare-improving trade cycle or claim chain in x^t ;
x^{t+1}	\leftarrow the result when σ^t is executed in x^t ;
t	$\leftarrow t + 1$;

end

OUTPUT: the deterministic assignment x^t ;

Algorithm 1: The v -Sequential Improvements Algorithm

Although [Proposition 8](#) is interesting by itself, the main reason for introducing it is to motivate the model of tinkering laid out in [Algorithm 1](#). Essentially, the algorithm starts with an arbitrary deterministic assignment and a valuation vector and then executes welfare-improving trade cycles and claim chains until no more exist.²⁰ This algorithm outputs a rank-efficient assignment in a finite number of steps. Formally,

Proposition 9. *Fix an ordinal market, \mathcal{M} . The v -Sequential Improvements Algorithm ([Algorithm 1](#)) terminates in a finite number of steps and outputs a rank-efficient assignment that is a member of $RV_v(\mathcal{M})$.*

The practical intuition of this tinkering algorithm is striking: given a value for 1st-choice assignments, 2nd-choice assignments, and so on, policymakers need only look for local welfare improvements (i.e., follow a greedy algorithm) to find a rank-efficient assignment.²¹ This sort of improvement process is natural, and in fact, it reflects what Al Roth and I found TFA doing when we first observed their process for matching teachers to regions. (See [footnote 2](#).) Administrators would spend about a week to match a batch of around 2,000 teachers, looking for non-Pareto-improving trade cycles where the good was perceived to outweigh the bad, periodically checking progress by calculating the rank distribution of the candidate assignment.

To be clear, [Algorithm 1](#) is meant to be descriptive and not prescriptive (i.e., it is a model of what tinkering policymakers are already doing, not

rank-inefficient assignment whose rank dominator can only be reached by executing two disjoint trade cycles. One transforms three 2nd choices into two 1st choices and a 3rd choice, while the other transforms a 1st and a 3rd choice into two 2nd choices. Neither cycle stochastically improves the rank distribution by itself, but on net, they transform one 2nd choice into a 1st choice, which is a clear stochastic improvement.

20. I assume that the starting assignment is deterministic to ensure that the algorithm terminates in finite time. This also holds when the starting assignment is nondeterministic, but constrained to have all entries rational.

21. In addition, note these local welfare improvements can be found with polynomial-time algorithms (e.g., the Bellman–Ford algorithm, described in [Leiserson et al. 2001](#)).

an algorithm to be suggested to them). That said, it is worth mentioning that its cycle-finding approach has parallels in other parts of the matching literature. The most obvious is the algorithm implied by the cyclic characterization of ordinal efficiency (Bogomolnaia and Moulin 2001; Katta and Sethuraman 2006). In addition, Erdil and Ergin (2008) look to “stable improvement cycles” as a way to improve efficiency when the assignment must be stable. Finally, in the operations-research literature, Klein (1967) uses a cycle-finding approach to solve minimum-cost flow problems.

5.2 Leeway for Improvement and the Desire to Tinker at HBS

TINKERING might not be a first-order issue if there isn’t much leeway to improve the originally selected match. In this subsection, I present an empirical setting where the leeway for improvement on a random serial dictatorship mechanism is large—so large that policymakers had to be restrained from the sort of tinkering just described. This suggests that market designers should be aware of the tinkering temptation faced by policymakers.

At HBS, first-year MBAs must participate in a global immersion program, known as Field Immersion Experiences for Leadership Development, or FIELD. They are assigned to a foreign company and remotely work on a project with that company during their first semester. FIELD culminates in a two-week trip over the winter break during which the MBAs present their work in person and make foreign business contacts.

In 2011 (the first year of the FIELD program), the match of students to countries was determined by an ordinal assignment mechanism that I designed jointly with Al Roth. At the beginning of their first semester, 900 MBAs ranked the 11 different countries to which they could be assigned. Once an MBA was matched to a country, the company she got within that country was administratively assigned without her input. The rankings submitted by the MBAs form the raw data for the empirical exercise that follows.

The mechanism I helped to design is a version of random serial dictatorship adapted to deal with constraints like demographic diversity; I describe it more fully in [Appendix B](#). For the purposes of this section, two features of the design are important. First, the mechanism is strategyproof, which means it is reasonable to take the submitted rankings as truthful (cf. Li 2017). Second, the MBAs were allowed to express indifferences in their submitted rankings. To keep the exercise in line with the model used in this paper, I randomly break those indifferences; however, the results are qualitatively similar if I don’t.

Essentially, the empirical exercise is to compare the rank distribution achieved by random serial dictatorship to the rank distribution achieved by a rank-value mechanism. Determining the leeway for improving the rank distribution in the FIELD setting provides evidence of the potential tinkering temptation faced by policymakers more broadly. In this exercise, I will need to choose valuation vectors to input into the rank-value mechanism.

Specifically, I consider a one-parameter family of valuation vectors where the valuation of a k th-choice allocation is given by

$$v_k(\alpha) = \left(\frac{11-k}{10} \right)^\alpha. \quad (4)$$

I call the strictly positive parameter, α , the **exponent** of the valuation. Regardless of the exponent, the valuation of a 1st-choice allocation is 1 and the valuation of an 11th-choice allocation is 0. For an exponent that is greater than one, the marginal value of moving from a $(k+1)$ st rank to a k th rank is decreasing in k , that is, it is more valuable to move agents from highly ranked allocations to even more highly ranked allocations than to rescue them from poor allocations.²² For an exponent less than one, this is reversed. When the exponent equals one, the value of moving from a $(k+1)$ st rank to a k th becomes independent of k , which means that the rank-value mechanism simply minimizes the average rank across all agents. The valuation vectors for the specific exponents I use in the empirical exercise (i.e., 1/10, 1, and 10) are plotted in [Figure B.1](#) of [Appendix B](#).

The empirical exercise is simple. First, I bootstrap 10,000 markets, where the randomness comes in how indifferences are broken and which agent ordering is used for serial dictatorship.²³ Then, I run random serial dictatorship and the rank-value mechanism for the three exponents listed in the previous paragraph. Finally, I compile all bootstraps to yield an expected rank distribution for each mechanism. These rank distributions are listed in [Figure 3](#).

The rank-value mechanisms give rank distributions that look markedly better than the rank distribution given by random serial dictatorship. The expected number of students who get their first or second choice increases from about 73 percentage points under random serial dictatorship to about 86 percentage points under the rank-value mechanisms. This is a roughly 18 percent increase, which corresponds to about 160 students of the 900 who participated.²⁴

The results from the empirical exercise suggest that tinkering might be tempting in other settings, which speaks to the relevance of the formal model in [Section 5.1](#). That said, if the tinkering process itself changes the rankings that agents submit to a mechanism, then policymakers could be

22. Mathematically, this follows from the derivative of [Equation 3](#) with respect to k .

23. More precisely, in each bootstrap, each indifference class in each MBA's ranking is replaced by a strict ordering of its component countries that is independently drawn from the uniform distribution over such strict orderings. The agent ordering used by the serial dictatorship mechanism is independently drawn from the uniform distribution over all agent orderings.

24. It is worth noting that doing the same empirical exercise with the probabilistic serial mechanism yields a rank distribution that is indistinguishable (on the scale of [Figure 3](#)) from that yielded by random serial dictatorship. Che and Kojima (2010) shows that this is asymptotically expected.

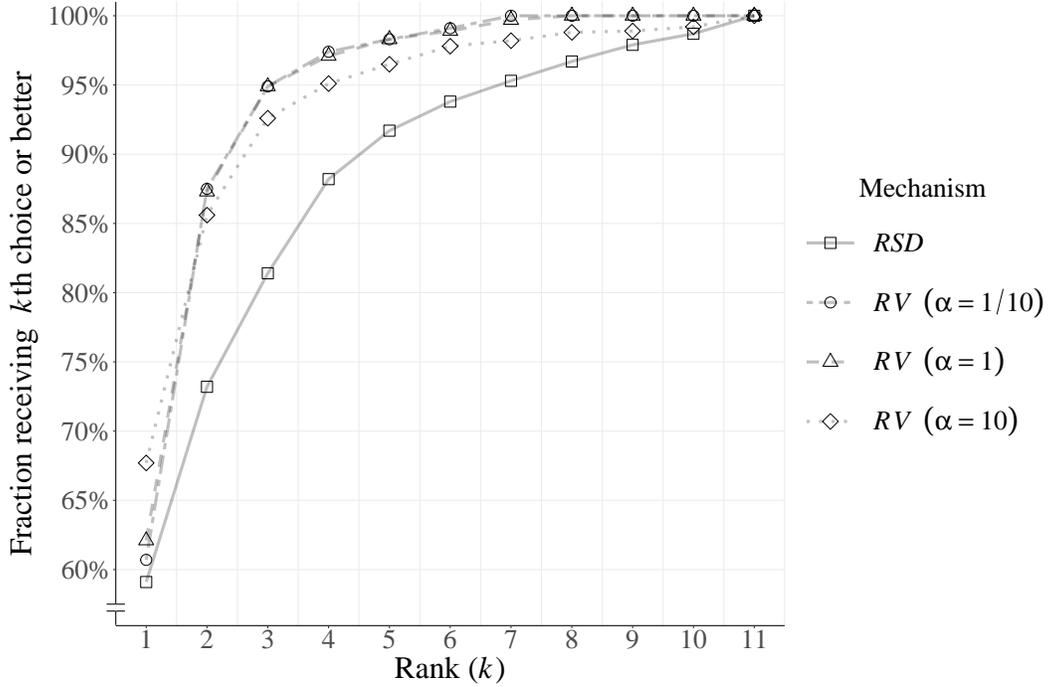


Figure 3: Average Rank Distributions for Random Serial Dictatorship and the Rank-Value Counterfactuals

NOTES: *RSD* stands for the random serial dictatorship mechanism and *RV* stands for the rank-value mechanism. The parenthetical numbers on the *RV* legend entries indicate the exponent (α) used by the valuation formula listed in Equation 4. The lines connecting the points are solely meant to help the reader visualize the rank distribution. All points overlap at rank 11, since all MBAs get an 11th-choice or better in the HBS setting.

chasing gains relative to reported preferences that deviate significantly from the truth. I turn to this concern next.

6 Incentives and Ranking Revelation

I BEGIN this section by showing that no ordinal assignment mechanism is both rank efficient and strategyproof. For administrators at HBS (see footnote 1), this impossibility result provided an important insight: when choosing between random serial dictatorship and a rank-value mechanism, HBS rejected the latter due to its lack of strategyproofness. In fact, the result (combined with insights from Section 5) even dissuaded HBS from tinkering ex post. Given the two-year MBA program, it seems plausible that an “overlapping generations” learning process might quickly converge.

This concern about lack of strategyproofness was not shared by administrators at TFA (see footnote 2): in the spring of 2011, they institutionalized and automated their tinkering approach by implementing a rank-value mechanism via linear programming. Administrators reasoned that teachers

only apply once to TFA, are geographically separated, and know little about what regions are popular relative to capacity. Note that the impossibility result does not imply that TFA’s reasoning is wrong: it merely shows that when facing a rank-efficient mechanism, an agent *can* have beliefs that lead her to deviate from truth-telling. It does not, however, show that she *must* do better to manipulate her report, regardless of beliefs (see, e.g., Roth and Rothblum 1999, Ehlers 2008, Featherstone and Niederle 2016, and Featherstone, Mayefsky, and Sullivan 2019).

After showing the impossibility result, I turn to modeling beliefs that make sense of TFA’s rationale. Ultimately, I find that whenever agents are asked to rank all object types (as they are in the TFA match), truth-telling is a best-response to beliefs meant as a stylized model of a general lack of information (see Section 6.2). When the agent is allowed to express which options are acceptable (i.e., preferred to the outside option), this result fails, even with so little information available to agents. However, an agent’s best-response to such beliefs can still be tightly characterized. Surprisingly, the form of the best-response depends on a seemingly insignificant design detail: what ranking the mechanism attaches to leaving an agent unmatched. In other words, it matters whether policymakers consider leaving an agent with her outside option as a distinctly bad outcome or as one potentially as good as being assigned.

6.1 Incompatibility of Rank Efficiency and Strategyproofness

AN ORDINAL assignment mechanism is **rank efficient** if, for any ordinal market input, it outputs an assignment that is rank efficient relative to that input. Of course, this guarantee of rank efficiency is less valuable without some concurrent assurance of truthful preference revelation. To examine whether such assurance exists, I model the ranking revelation problem. Ultimately, I will find that agents can sometimes gain from manipulating their preference report. In other words, when facing a rank-efficient mechanism, truth-telling isn’t a one-size-fits-all best-response.

I assume an agent knows with certainty the set of object types and her true ranking over those types. She does not know (with certainty, at least) how the rest of the market will look. Define a **market realization** to be an ordered quadruple consisting of an ordinal assignment mechanism, a set of other agents, a profile of the other agents’ submitted rankings, and a capacity profile. Given a set of object types, \mathcal{S} , a market realization is **conformable to \mathcal{S}** (or \mathcal{S} -conformable, for short), if the rankings and capacity profile in the market realization are over \mathcal{S} .

Given her knowledge of the set of object types, \mathcal{O} , if an agent, a , thinks the \mathcal{O} -conformable market realization is $\mathcal{M}_{-a} = (\Psi, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$, then when she reports a ranking, \tilde{r}_a , she also thinks the ordinal market reported to Ψ is $\mathcal{M} = (\mathcal{A}_{-a} \cup \{a\}, \mathcal{O}, (\tilde{r}_a; \tilde{r}_{-a}), q)$.²⁵ Hence, she thinks the final assignment

25. Per standard game theory notation, \tilde{r}_{-a} represents the rank profile of all agents except for

will be $\Psi(\mathcal{M})$.

An ordinal assignment mechanism, Ψ , is **strategyproof** if for every set of object types, \mathcal{O} , every true ranking that an agent, a , could have, r_a , and every \mathcal{O} -conformable market realization with Ψ as its mechanism, $(\Psi, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$, agent a weakly prefers (in the first-order stochastic sense) her allocation when she reports her true ranking, r_a , to her allocation when she reports any other ranking, \tilde{r}_a . That is, if $x = \Psi(\mathcal{A}_{-a} \cup \{a\}, \mathcal{O}, (r_a; \tilde{r}_{-a}), q)$ and $\tilde{x} = \Psi(\mathcal{A}_{-a} \cup \{a\}, \mathcal{O}, (\tilde{r}_a; \tilde{r}_{-a}), q)$, then $N^x(k; a) \geq N^{\tilde{x}}(k; a)$ for all ranks, k .

Intuitively, strategyproofness is a guarantee that an agent can't be hurt by reporting her true ranking (i.e., **truth-telling**), regardless of her underlying cardinal utility. Unfortunately, regardless of the ordinal assignment mechanism, strategyproofness cannot coexist with rank efficiency—a result that kept HBS from adopting a rank-value mechanism. Formally,

Proposition 10. *No ordinal assignment mechanism is both rank efficient and strategyproof.*

This proposition is closely related to Theorem 2 of Bogomolnaia and Moulin (2001), which shows that no ordinal assignment mechanism is ordinally efficient, strategyproof, and *satisfies equal treatment of equals* (i.e., agents with identical reported rankings get identical allocations).²⁶ Note that equal treatment of equals is required for the Bogomolnaia and Moulin theorem to hold, since a deterministic serial dictatorship is both strategyproof and ordinally efficient. (Proposition 1(c) and Proposition 4 combine to establish the latter property.) By strengthening ordinal efficiency to rank efficiency, Proposition 10 is able to dispense with the final condition of the Bogomolnaia and Moulin theorem.

The intuition for Proposition 10 is captured by the example in Figure 2. Someone has to get object o_4 , and agent a_4 ranks it 3rd, while the other agents rank it 4th. Rank efficiency then dictates that agent a_4 gets stuck with the bad object. But, if agent a_4 submits the false ranking that switches objects o_3 and o_4 , she looks like she dislikes object o_4 as much as everyone else, allowing a rank-efficient mechanism to give her objects she prefers.

6.2 Modeling a Lack of Information

WHEN FACING a strategyproof mechanism, an agent will weakly prefer truth-telling, regardless of her beliefs about the probabilities that given market realizations occur. Hence, the incompatibility of rank efficiency and strategyproofness means that, when facing a rank-efficient mechanism, an agent *can* have beliefs that lead her to manipulate her preference reports.

a , while $(\tilde{r}_a; \tilde{r}_{-a})$ represents the profile where the ranking of agent a is \tilde{r}_a and the rankings of all other agents are given by \tilde{r}_{-a} .

26. Proposition 10 is also similar to Theorem 1 of Zhou (1990), although the latter result concerns a setting where agents report their von Neumann–Morgenstern utilities directly.

It does not, however, mean she *must* do better to manipulate her report, regardless of her beliefs.

This insight is crucial to making any sense of how TFA could justify the adoption of a rank-value mechanism. For the rest of [Section 6](#), I model agents with a general lack of information about supply and demand. Such agents can best-respond with a simple class of revelation strategies, and in fact, in environments without outside options, such agents can best-respond by truth-telling.

Before showing this, however, I must lay some groundwork. First, I introduce the method for selecting an assignment when the arg max of a welfare maximization LP ([Equation 1](#)) fails to be a singleton (see [Section 4.2](#)). Then, I introduce my model of agent beliefs when agents face a rank-value mechanism and have little information. Finally, I define a few classes of ranking strategies that will feature prominently in [Section 6.3](#).

6.2.1 Tiebreakers for Welfare Maximization Mechanisms

By definition, all assignments in the arg max of the welfare maximization LP attain the same objective value. The task is to break this tie, which I will do by taking advantage of an important fact: the arg max of the welfare maximization LP is the convex hull of a set of deterministic assignments.²⁷

Given this fact, a natural tiebreaking approach is to take that set of deterministic assignments, order the agents, and run an analogue of serial dictatorship. When an agent’s turn comes, she finds her most preferred object of those she gets in one of the remaining assignments and then eliminates all assignments that give her anything worse. Obviously, only one deterministic assignment will remain at the end of this procedure, since every agent has selected exactly one object type for herself.

Formally, fix an ordinal market, $(\mathcal{A}, \mathcal{O}, r, q)$, let \mathcal{X} be a convex hull of deterministic assignments, and let π be a deterministic agent ordering. Then, the **tiebreaker over \mathcal{X} given π** , denoted $b_\pi(\mathcal{X})$, returns the output of [Algorithm 2](#). This definition is easily extended to a random agent ordering, Π : the **tiebreaker over \mathcal{X} given Π** is the expectation of the random matrix that results from the tiebreaker given the deterministic realization of Π .

So far, I have argued that the tiebreaker procedure is a valid selection rule for non-singleton arg maxes of the welfare maximization LP. However, one might worry that some elements of the arg max cannot be reached by any tiebreaker. This concern is ultimately unfounded: any element of the arg max can be reached by the tiebreaker given some random agent ordering.²⁸

27. Say it weren’t. Then, there would be some assignment, x , in the arg max that is not the convex combination of deterministic assignments in the arg max. But from Budish et al. (2013), x is the convex combination of some set of deterministic assignments. Hence one of those assignments, which I’ll call y , is not in the arg max. But then, constructing a new assignment by dropping y from the support of x would yield a new assignment with a greater welfare objective than x , contradicting that x was in the arg max in the first place. See [Proposition A.2\(a\)](#) in [Appendix A](#).

28. Since any assignment in the arg max is ex post efficient, [Proposition 4](#) gives a random

INPUT: a convex hull of deterministic assignments, \mathcal{X} , and a deterministic agent ordering, π ;

$\mathcal{X}_0 \leftarrow$ the set of deterministic assignments in \mathcal{X} ;

for $t \leftarrow 1$ to $|\mathcal{A}|$ do

$\mathcal{O}_t \leftarrow$ the set of object types that agent $\pi(t)$ gets from at least one assignment in \mathcal{X}_{t-1} ;

$o_t^* \leftarrow$ the object type in \mathcal{O}_t that agent $\pi(t)$ most prefers;

$\mathcal{X}_t \leftarrow$ the set of assignments in \mathcal{X}_{t-1} that give object type o_t^* to agent $\pi(t)$;

end

OUTPUT: the sole element in $\mathcal{X}_{|\mathcal{A}|}$;

Algorithm 2: The Tiebreaker Algorithm

6.2.2 Modeling Beliefs

An agent’s **beliefs** are encoded by a probability distribution over market realizations. Given beliefs, the **expected allocation** for an agent, a , from a reported ranking, \tilde{r}_a , is her allocation in the assignment $\mathbb{E}[\Psi(\{a\} \cup \mathcal{A}_{-a}, \mathcal{O}, (\tilde{r}_a; \tilde{r}_{-a}), q)]$, where the expectation is over the mechanism, Ψ , the set of other agents, \mathcal{A}_{-a} , the submitted rankings of those agents, \tilde{r}_{-a} , and the capacity vector, q . Given this structure, the model of beliefs that I use in [Section 6.3](#) has two main components.

Facing a Rank-Value Mechanism. Intuitively, the agent should know the ordinal assignment mechanism is a rank-value mechanism that uses the tiebreaker defined by [Algorithm 2](#). However, arbitrary beliefs should be allowed concerning the valuation vector, v , and the deterministic ordering of agents that will be used by the tiebreaker, π . That is, the agent might know these things with certainty, but she also might not. Formally, I say an agent is **facing a rank-value mechanism** when her beliefs place zero probability on any market realization whose ordinal assignment mechanism isn’t of the form $b_\pi \circ RV_v$, for some valuation, v , and deterministic ordering, π .²⁹ Note that the specific valuation and deterministic ordering is allowed to vary across market realizations. Formulating beliefs in terms of market realizations with a deterministic tiebreaker is without loss of generality, since random tiebreakers are defined as lotteries over deterministic tiebreakers.

Low-Information Beliefs. I model an agent’s lack of information by placing constraints on beliefs that rely on a swapping operation first introduced

ordering that can generate the assignment from serial dictatorship. The tiebreaker given this random ordering will also yield the assignment. See [Proposition A.2\(c\)](#) in [Appendix A](#).

29. Following standard notation, $b_\pi \circ RV_v$ represents the composition that first applies the v -rank-value mechanism and then applies the tiebreaker of [Section 6.2.1](#), given π .

in Roth and Rothblum (1999). For some profile of rankings, r , and two object types, o and o' , let $r^{o \leftrightarrow o'}$ denote the profile where all agents switch their rankings of o and o' , that is, $r^{o \leftrightarrow o'}$ is the same as r except for all agents, a , we have $r_{ao}^{o \leftrightarrow o'} = r_{ao'}$ and $r_{ao'}^{o \leftrightarrow o'} = r_{ao}$. The same definition can be used, *mutatis mutandis*, to apply the swapping operation to an individual agent's ranking, to a profile of rankings excluding one agent, and to a capacity profile.

Now, I can introduce the constraints on beliefs. For an agent, a , who knows the set of object types is \mathcal{O} , for two object types, o and o' , I call her beliefs **$\{o, o'\}$ -symmetric** if the probability of any \mathcal{O} -conformable market realization, $(\Psi, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$, is equal to the probability of $(\Psi, \mathcal{A}_{-a}, \tilde{r}_{-a}^{o \leftrightarrow o'}, q^{o \leftrightarrow o'})$. Similarly, for a subset of object types, \mathcal{S} , the beliefs of an agent are **\mathcal{S} -symmetric** if they are $\{o, o'\}$ -symmetric for any o and o' in \mathcal{S} .

Intuitively, belief symmetry models how agents deal with a lack of information about the supply or demand of one object type relative to another. More broadly, one might interpret symmetric beliefs as a tractable model of an agent's inability to compute the relative probabilities of potential welfare-improving cycles. This perspective is further discussed in [Section 6.4](#).

6.2.3 Truthfulness, Extension, and Truncation

Finally, I define a few important classes of ranking-revelation strategies. Consider an agent, a , who knows the set of object types is \mathcal{O} , and let \mathcal{S} be a subset of \mathcal{O} . A reported ranking from agent a is **truthful about \mathcal{S}** when for any object types in \mathcal{S} , if one is ranked higher than the other in her true ranking, then it is also ranked higher in the report. That is, a report, \tilde{r}_a , is truthful about \mathcal{S} relative to the true ranking, r_a , if, for every pair of object types, o and o' , in \mathcal{S} , when $r_{ao'} < r_{ao}$, it is also true that $\tilde{r}_{ao'} < \tilde{r}_{ao}$. **Truth-telling** is the strategy that is truthful about \mathcal{O} .

A submitted ranking that is truthful about all non-null object types, but declares some truly unacceptable object types to be acceptable is an **extension**. Conversely, a submitted ranking that is truthful about all non-null object types, but declares some truly acceptable object types to be unacceptable is a **truncation**. Along with truth-telling, both of these types of revelation strategies play an important role for agents with symmetric beliefs in settings with outside options. In the matching setting, truncations have been seen before; however, extensions are (to the best of my knowledge) new.

Of course, if an agent submits an extension of her true ranking, it introduces the possibility that she is assigned to an object type that is unacceptable under her true ranking. This raises the question of what utility she should attach to such an outcome. Since the null object type is meant to represent an outside option, it makes sense that she would simply refuse the assignment and take her outside option. To model this, I assume that when an agent is assigned an object type she finds truly unacceptable, she treats it the same as if she had been assigned the null object type. Under this assumption, a specific extension becomes important. The **full exten-**

sion is the extension that places all non-null object types in their truthful order, while ranking the null object type last.

6.3 Ranking Revelation with Low-Information Beliefs

Now, I can discuss how agents best-respond to beliefs of the form described in [Section 6.2](#). Most of the results derive from a simple observation: symmetric beliefs about a set of object types guarantee that an agent believes she can't be hurt by being truthful about that set. Formally,

Proposition 11. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let S be a subset of \mathcal{O} , and let the agent have S -symmetric beliefs.*

Then, for any report that isn't truthful about S , there exists another report that is truthful about S such that the agent weakly prefers (in the first-order stochastic sense) her expected allocation from the latter report to her expected allocation from the former.

Intuitively, consider a market realization where when the agent ranks two object types, o and o' , in their true order, she gets another object type, o'' . When beliefs are symmetric, such realizations are just as likely as realizations where agent a gets o'' from switching her reported rankings of o and o' . The advantage of truthfully ordering the object types comes from market realizations in which, of o and o' , the agent gets the object type she ranks higher.

Using [Proposition 11](#) as a starting point, I will now consider three scenarios in which best-responses can be tightly characterized.

6.3.1 Truth-Telling When Outside Options Aren't Ranked

[Proposition 11](#) clearly implies that when an agent has symmetric beliefs about all object types, she believes she can't be hurt by truth-telling. Formally,

Proposition 12. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let her have \mathcal{O} -symmetric beliefs.*

Then, the agent weakly prefers (in the first-order stochastic sense) her expected allocation from truth-telling to her expected allocation from any other report.

But is it sensible that an agent should have \mathcal{O} -symmetric beliefs? If the null object type, \emptyset , is in \mathcal{O} , then \mathcal{O} -symmetry effectively imposes the belief that non-null object types are just as likely to be unacceptable as they are to be acceptable. [Proposition 12](#) seems best applied when \emptyset is not in \mathcal{O} , that is, when outside options are not ranked.

Two situations fit this bill. One is when there is no plausible outside option, as in military branch-of-service matching ([Sönmez and Switzer 2013](#)). The other is when agents are simply not allowed to explicitly rank their outside options. This is quite common in school assignment, where students

receive an assignment to *some* school, regardless of their submitted ranking. In these settings, failing to rank all schools is effectively equivalent to having a bottom indifference class consisting of all unranked schools. If the school system can't give a student a school she ranks, it assigns her to an unranked school that has leftover capacity. TFA also doesn't allow its teachers to explicitly rank their outside options. This policy is rooted in the organization's faith in its ability to persuade: it likes the opportunity to change a teacher's mind about accepting an assignment to a region that she might have perceived as unacceptable at the time of rank submission.

6.3.2 Extension When Outside Options Are Ranked

When agents are allowed to rank their outside options, i.e., when \emptyset is in \mathcal{O} , a more sensible "low-information" assumption is that beliefs are $(\mathcal{O} \setminus \{\emptyset\})$ -symmetric. With such beliefs, an agent still believes she does weakly better to submit a ranking that is truthful about the non-null object types. However, she might misrepresent where the null object type lies. Specifically,

Proposition 13. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let \mathcal{O} include the null object type, \emptyset , and let the agent have $(\mathcal{O} \setminus \{\emptyset\})$ -symmetric beliefs.*

Then, the agent weakly prefers (in the first-order stochastic sense) her expected allocation from reporting the full extension to her expected allocation from any other report.

Essentially, the proposition does two things: it rules out truncation and insists on full extension. Why is truncation ruled out? If an agent ranks \emptyset second when it is truly third, then the v -rank-value mechanism gets v_2 points for giving her \emptyset , without incurring any opportunity cost (since \emptyset is non-scarce). If the agent instead pushes \emptyset to the third position, it becomes more costly for the mechanism to give her \emptyset . Why full extension? A parallel argument shows that when an agent reports an extension, she penalizes the mechanism more for giving her \emptyset , but at the cost of potentially being assigned an unacceptable object type. However, this isn't actually a downside, since I assume she can costlessly refuse that assignment.

Practically, [Proposition 13](#) supports a common practice among school assignment mechanisms: not allowing students to rank their outside options. If they are allowed to rank their outside options, they can actually be hurt by truth-telling. Preventing this by omitting the outside option from the set of schools to be ranked can be construed as one way to "level the playing field" between naive and sophisticated players ([Pathak and Sönmez 2008](#)).

6.3.3 Truncation When Outside Options Are Ranked

In [Section 6.3.2](#), the incentive to submit an extension was driven by the mechanism treating the ranking of null and non-null object types the same. Policymakers, however, might feel like giving an agent her outside option is, in some sense, worse than other outcomes. School districts often report

how many students got a first choice, second choice, and so on, but also report how many were not matched to one of their ranked choices. (See the publicly released reports cited at the beginning of the introduction for examples from New York City and San Francisco schools.)

Formally, this is captured by how submitted preferences are converted to rankings. Consider eliciting a strict ordinal preference over the set of object types and then using a **rank function** to convert that preference into a ranking to be used by the rank-value mechanism. The rank function I have tacitly used throughout the paper assigns a rank to an object type, o , according to the function $\rho_o(>_a) = 1 + |\{o' : o' >_a o\}|$, that is, an object type's rank is one more than the number of object types preferred to it. While this is sensible, alternatives are conceivable. In this subsection, I will briefly consider the rank function

$$\rho_o(>_a) = \begin{cases} 1 + |\{o' : o' >_a o\}| & \text{if } o >_a \emptyset \\ |\mathcal{O}| & \text{if } o = \emptyset \\ |\mathcal{O}| + 1 + |\{o' : \emptyset >_a o' >_a o\}| & \text{if } \emptyset >_a o. \end{cases} \quad (5)$$

It is meant to capture the intuition that giving an agent the null object type is, in some sense, worse than any other acceptable outcome. Acceptable object types are ranked as before, but starting at \emptyset , the rank jumps to $|\mathcal{O}|$ and then increases from there for unacceptable object types. Now, \emptyset is always ranked $|\mathcal{O}|$ th, regardless of where it is in the submitted preference.

With very few minor changes, the results from Sections 3–5 continue to hold when strict preferences are mapped to rankings according to Equation 5, since it is just a different way to aggregate welfare. However, the new formula does change strategic incentives. Before, agents were incentivized to extend because it allowed them to penalize the mechanism more for giving them \emptyset . With the new formula, this is unnecessary; in fact, the new rank function gives agents an incentive to *truncate*. Intuitively, if an agent refuses to allow the mechanism to give her a low-ranked, but truly acceptable, object type, then the mechanism must either give her a better object type or pay the price for giving her \emptyset . Depending on which the mechanism chooses, truncation could prove profitable to the agent.

Formally, consider an agent, a , who knows that the set of object types is \mathcal{O} . Given an \mathcal{O} -conformable market realization, $(\Psi, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$, when she submits a strict preference, $>_a$, let her expect the assignment to be $\Psi(\mathcal{A}_{-a} \cup \{a\}, \mathcal{O}, (\rho(>_a); \tilde{r}_{-a}), q)$, where the rank function ρ is given by Equation 5. Given this setup, the result can be stated as follows.

Proposition 14. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Let \mathcal{O} contain \emptyset , and let the agent have $(\mathcal{O} \setminus \{\emptyset\})$ -symmetric beliefs. Further, assume when the agent submits a strict preference, the ranking submitted on her behalf will be computed using the rank function ρ , as defined in Equation 5.*

Then, for any reported strict preference that isn't truncation or truth-telling, there exists another reported strict preference that is truncation or truth-telling such that the agent weakly prefers (in the first-order stochastic sense) her expected allocation from the latter report to her expected allocation from the former.

6.4 Discussion of Ranking Revelation Results

THAT NO no rank-efficient mechanism is strategyproof ([Proposition 10](#)) implies that when facing a rank-efficient mechanism, there exist beliefs to which an agent best-responds by deviating from truth-telling. Of course, this doesn't preclude the existence of beliefs to which an agent can best-respond by truth-telling. Sections [6.2](#) and [6.3](#) explore one such class of beliefs in an attempt to rationalize TFA's adoption of a rank-value mechanism.

When outside options are not an issue, an agent can respond to symmetric beliefs by truth-telling ([Proposition 12](#)). Of course, taken literally, belief symmetry is a strong condition. Alternatively, one might interpret belief symmetry as a tractable model of an agent's inability to compute the relative probabilities of potential welfare-improving cycles. Such an inability could be rooted in something outside of the model, such as lack of information about the details of the algorithm. In the words of Roth and Rothblum ([1999](#)): "...[when] simply trying to explain [agent] behavior, we may want to consider the assumption of symmetric information not merely literally, but also as a metaphor for a wider class of uncertainty." Rubinstein ([1991](#)) and Rubinstein ([2006](#)) consider such a perspective more broadly in the context of economic theory.

When there are outside options, however, [Propositions 13](#) and [14](#) demonstrate that even with very little information, agents may have the incentive to deviate from truth-telling (a result reminiscent of Roth and Rothblum [1999](#)). It is worth emphasizing that the difference between the two propositions hinges on an otherwise unremarkable detail of the design: the rank given to the object type that represents agents' outside options. Understanding which stance a tinkering policymaker takes can shed light on how agents respond. Clearinghouses in the field tend to report unmatched agents as a distinct category, which speaks to the approach of [Proposition 14](#); however, the approach of [Proposition 13](#) resonates with the idea that it is better to leave an agent unmatched than to give her an assignment that she will refuse ex post. One thing is clear: as usual in market design, the details matter (Roth [2002](#); Duflo [2017](#)).

7 Conclusion

THE RANK DISTRIBUTION is a natural summary statistic for the quality of a match, and if policymakers tinker by implementing local improvements to the welfare objective, they will find a rank-efficient assignment. The apparent gains from doing so can be large. As such, a major contribution

of this paper is modeling what could be a common practice in the field.

However, as market designers, the incompatibility of rank efficiency and strategyproofness should make us reluctant to recommend rank-value mechanisms to policymakers. Still, the results from [Section 5.2](#) are tantalizing: they suggest that if the incentive problem could somehow be circumvented, the efficiency gains would be substantial. This situation is not unique; in fact, there is a growing literature highlighting the costs of strategyproofness in settings where impossibility theorems preclude its co-existence with a pertinent efficiency concept. For instance, in the context of stable matching, both Erdil and Ergin ([2008](#)) and Abdulkadiroğlu, Pathak, and Roth ([2009](#)) look at how much efficiency is lost when indifference in school priorities are broken poorly. Correcting these “mistakes”, however, undermines strategyproofness.

To be clear, lack of strategyproofness should not immediately rule out a mechanism. In the words of Day and Milgrom ([2008](#)),

...it is customary in mechanism design theory to impose incentive constraints first, investigating other aspects of performance only later. It is, of course, equally valid to begin with other constraints, and such an approach can be useful. To the extent that optimization is only an approximation to the correct behavioral theory for [agents], it is interesting to investigate how closely incentive constraints can be approximated when other constraints are imposed first.

A lack of strategyproofness simply shows that there will be incentive problems in at least *some* environments. It does not, however, rule out the existence of environments in which these problems are mitigated.³⁰

Is a rank-value mechanism ever a good design choice? The stylized model of beliefs laid out in [Section 6.2](#) attempts to provide an answer for environments in which agents have little information. In such environments, when outside options are ranked, [Propositions 13](#) and [14](#) suggest that deviation from truth-telling should still be a concern. However, when outside options aren't ranked, [Proposition 12](#) shows that truth-telling is a best response, providing at least some theoretical rationale for TFA's decision to implement a rank-value mechanism.

While there are many papers on the costs of strategyproofness, this paper also considers a lesser-known concern. As a principle, Pareto agnosticism—that is, insisting the only way to improve an assignment is to make all agents better off—can come with its own costs. If the policymaker truly considers all k th-choice allocations as interchangeable,

30. In fact, there are many market-design papers that focus on non-strategyproof mechanisms whose incentive properties improve in environments that are, in some sense, large. See Roth and Peranson ([1999](#)), Kojima and Pathak ([2009](#)), Che and Kojima ([2010](#)), Kojima and Manea ([2010](#)), Budish ([2011](#)), Kojima, Pathak, and Roth ([2013](#)), and Azevedo and Budish ([2018](#)), among others.

then (barring incentive concerns) by allowing her to express this, the rank-value mechanism can yield big efficiency gains over mechanisms that are merely Pareto efficient.

Concerning Pareto agnosticism, the lessons from this paper are three-fold. First, although economists might not be in a position to make ethical judgments (in the sense of Harsanyi 1955) about interpersonal utility comparison, policymakers usually are. It then makes sense for market designers to use mechanisms that allow policymaker judgments to be transparently inputted. The valuation plays this role in the rank-value mechanism. Second, by explicitly arming policymakers with the power to input such judgments, welfare can be significantly improved if incentive problems can be overcome. And finally, market designers should be concerned about tinkering when policymakers are given mechanisms that do not achieve their implicit goals. Strategyproofness for the agents does not matter unless the mechanism is also, in some sense, “strategyproof” for the policymaker. Practically, market designers must be vigilant to ensure that policymakers are convinced (and remain convinced) that the goals of a mechanism are their goals as well. Such “plumbing” details matter (Duflo 2017).

The fact that there are rank-efficient mechanisms in the field could mean that they are successful in some situations. Alternatively, it might mean that well-intentioned policymakers, ignorant of incentive concerns, are implementing bad policies. Market designers need to know whether they should be correcting the impulse to depart from strategyproofness or sometimes suggesting that it is a small price to pay for potentially big welfare gains. Which advice is good very likely depends on the environment in which the mechanism is meant to serve. Understanding when (if ever) market designers can comfortably suggest a rank-value mechanism is an important open question.

References

- 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”. *American Economic Review* 99 (5): 1954–78.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. “Ordinal Efficiency and Dominated Sets of Assignments”. *Journal of Economic Theory* 112 (1): 157–172.
- Abdulkadiroğlu, Atila, and Tayfun Sönmez. 1998. “Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems”. *Econometrica*: 689–701.
- . 2003. “School Choice: A Mechanism Design Approach”. *American Economic Review* 93 (3): 729–747.
- Azevedo, Eduardo M., and Eric Budish. 2018. “Strategy-proofness in the Large”. *The Review of Economic Studies* 86 (1): 81–116.

- Birkhoff, Garrett. 1946. "Tres Observaciones Sobre el Algebra Lineal". *Universidad Nacional de Tucumán Revista, Serie 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.
- Bronfman, Slava, Avinatan Hassidim, Arnon Afek, Assaf Romm, Rony Shreberk, Ayal Hassidim, and Anda Massler. 2015. "Assigning Israeli Medical Graduates to Internships". *Israel Journal of Health Policy Research* 4 (1): 6.
- Budish, Eric. 2011. "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes". *Journal of Political Economy* 119 (6): 1061–1103.
- Budish, Eric, and Estelle Cantillon. 2012. "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard". *American Economic Review* 102 (5): 2237–2271.
- Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. "Designing Random Allocation Mechanisms: Theory and Applications". *American Economic Review* 103 (2): 585–623.
- Budish, Eric, and Judd B. Kessler. 2018. "Bringing Real Market Participants' Real Preferences into the Lab: An Experiment that Changed the Course Allocation Mechanism at Wharton". *NBER Working Paper*, no. w22448.
- Che, Yeon-Koo, and Fuhito Kojima. 2010. "Asymptotic Equivalence of Probabilistic Serial and Random Priority Mechanisms". *Econometrica* 78 (5): 1625–1672.
- Davis, Jonathan M.V., and Pablo Montagnes. 2020. "Designing Organizational Versus Public Markets". *Working paper*.
- Day, Robert, and Paul Milgrom. 2008. "Core-Selecting Package Auctions". *International Journal of Game Theory* 36 (3): 393–407.
- Duflo, Esther. 2017. "Richard T. Ely Lecture: The Economist as Plumber". *American Economic Review* 107 (5): 1–26.
- Ehlers, Lars. 2008. "Truncation Strategies in Matching Markets". *Mathematics of Operations Research* 33 (2): 327–335.
- Erdil, Aytek, and Haluk Ergin. 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice". *American Economic Review* 98 (3): 669–689.
- Featherstone, Clayton R., Eric S. Mayefsky, and Colin D. Sullivan. 2019. "Why Do Some Clearinghouses Yield Stable Outcomes: Experimental Evidence on Out-of-Equilibrium Truth-telling". *Working paper*.
- Featherstone, Clayton R., and Muriel Niederle. 2016. "Boston Versus Deferred Acceptance in an Interim Setting: An Experimental Investigation". *Games and Economic Behavior* 100:353–375.
- Harsanyi, John C. 1975. "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory". *American Political Science Review* 69 (2): 594–606.
- . 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility". *Journal of Political Economy* 63 (4): 309–321.

- . 1986. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. Cambridge University Press.
- Irving, Robert W., Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E. Paluch. 2006. “Rank-Maximal Matchings”. *ACM Transactions on Algorithms (TALG)* 2 (4): 602–610.
- Jackson, Matthew O. 2001. “A Crash Course in Implementation Theory”. *Social Choice and Welfare* 18 (4): 655–708.
- 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.
- Kavitha, Telikepalli, and Chintan D. Shah. 2006. “Efficient Algorithms for Weighted Rank-Maximal Matchings and Related Problems”. In *International Symposium on Algorithms and Computation*, 153–162. Springer.
- Klein, M. 1967. “A Primal Method for Minimal Cost Flows with Applications to the Assignment and Transportation Problems”. *Management Science*: 205–220.
- Kojima, Fuhito, and Mihai Manea. 2010. “Incentives in the Probabilistic Serial Mechanism”. *Journal of Economic Theory* 145 (1): 106–123.
- Kojima, Fuhito, and Parag A. Pathak. 2009. “Incentives and Stability in Large Two-Sided Matching Markets”. *American Economic Review* 99 (3): 608–27.
- Kojima, Fuhito, Parag A. Pathak, and Alvin E. Roth. 2013. “Matching with Couples: Stability and Incentives in Large Markets”. *Quarterly Journal of Economics* 128 (4): 1585–1632.
- Kojima, Fuhito, and M. Utku Ünver. 2014. “The “Boston” School-Choice Mechanism: An Axiomatic Approach”. *Economic Theory* 55 (3): 515–544.
- Leiserson, Charles Eric, Ronald L Rivest, Thomas H Cormen, and Clifford Stein. 2001. *Introduction to Algorithms*. Vol. 6. MIT Press, Cambridge, MA.
- Li, Shengwu. 2017. “Obviously Strategy-Proof Mechanisms”. *American Economic Review* 107 (11): 3257–87.
- Manea, Mihai. 2008. “A Constructive Proof of the Ordinal Efficiency Welfare Theorem”. *Journal of Economic Theory* 141 (1): 276–281.
- . 2009. “Asymptotic Ordinal Inefficiency of Random Serial Dictatorship”. *Theoretical Economics* 4 (2): 165–197.
- . 2007. “Serial Dictatorship and Pareto Optimality”. *Games and Economic Behavior* 61 (2): 316–330.
- McLennan, Andrew. 2002. “Ordinal Efficiency and the Polyhedral Separating Hyperplane Theorem”. *Journal of Economic Theory* 105 (2): 435–449.
- National Residency Matching Program. 2019. “Results and Data: 2019 Main Residency Match”. Visited on 06/26/2019. <http://www.sfusd.edu/en/enroll-in-sfusd-schools/monitoring-student-assignment/student-assignment-system.html>.
- New York City Department of Education. 2018. “Chancellor Carranza Announces Record High Percentage of Families Receiving Top Pre-K Choice”. *Press Release*

- from Chancellor's Press Office. Visited on 06/03/2019. <https://www.schools.nyc.gov/about-us/news/announcements/contentdetails/2018/05/08/chancellor-carranza-announces-record-high-percentage-of-families-receiving-top-pre-k-choice>.
- . 2009. "More Than 80 Percent of Students Admitted to a Top Choice High School for Fourth Consecutive Year". *High School Admissions Press Release*. Visited on 06/03/2019. https://web.archive.org/web/20090330075332/http://schools.nyc.gov/Offices/mediarelations/NewsandSpeeches/2008-2009/20090325_hs_admissions.htm.
 - Paluch, Katarzyna. 2013. "Capacitated Rank-Maximal Matchings". In *International Conference on Algorithms and Complexity*, 324–335. Springer.
 - Pathak, Parag A., and Tayfun Sönmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated players in the Boston Mechanism". *American Economic Review* 98 (4): 1636–52.
 - Pycia, Marek. 2019. "Evaluating with Statistics: Which Outcome Measures Differentiate Among Matching Mechanisms?" *Working paper*.
 - Rawls, John. 1972. *A Theory of Justice*. Oxford, Clarendon Press.
 - Rockafellar, R. Tyrrell. 1970. *Convex Analysis*. Vol. 28. Princeton University Press.
 - 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". *American Economic Review*: 415–440.
 - . 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics". *Econometrica* 70 (4): 1341–1378.
 - . 1984. "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory". *Journal of Political Economy*: 991–1016.
 - Roth, Alvin E., and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design". *American Economic Review* 89 (4): 748–780.
 - Roth, Alvin E., and Uriel G. Rothblum. 1999. "Truncation Strategies in Matching Markets—In Search of Advice for Participants". *Econometrica* 67 (1): 21–43.
 - Roth, Alvin E., and Marilda A. O. Sotomayor. 1992. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Vol. 18. Cambridge University Press.
 - Rubinstein, Ariel. 1991. "Comments on the Interpretation of Game Theory". *Econometrica* 59 (4): 909–24.
 - . 2006. "Dilemmas of an Economic Theorist". *Econometrica* 74 (4): 865–883.
 - San Francisco Unified School District. 2018a. "2018–19 March Placement Results (All Grade Levels)". Prepared by *Educational Placement Center*. Visited on 06/03/2019. http://www.sfusd.edu/en/assets/sfusd-staff/enroll/files/2018-19/2018-19_Placement_Results_TK-13.pdf.
 - . 2018b. "Enrollment Highlights 2018". Prepared by *Educational Placement Center*. Visited on 06/03/2019. <http://www.sfusd.edu/en/enroll-in-sfusd-schools/monitoring-student-assignment/student-assignment-system.html>.

- . 2011. “Student Assignment: March 2011 School Assignment Offers”. Prepared by *Educational Placement Center*. Visited on 06/03/2019. http://www.sfusd.edu/en/assets/sfusd-staff/enroll/files/Post_March/Update%20April%2013%202011_Revised.pdf.
- Sönmez, Tayfun, and Tobias B. Switzer. 2013. “Matching with (Branch-of-Choice) Contracts at the United States Military Academy”. *Econometrica* 81 (2): 451–488.
- Svensson, Lars-Gunnar. 1994. “Queue Allocation of Indivisible Goods”. *Social Choice and Welfare* 11 (4): 323–330.
- Thomson, William. 2011. “Fair Allocation Rules”. Chap. 21 in *Handbook of Social Choice and Welfare*, ed. by Kenneth J. Arrow, Amartya Sen, and Kotaro Suzumura, 2:393–506. Elsevier.
- 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.
- . 2005. “On the Survival of Some Unstable Two-Sided Matching Mechanisms”. *International Journal of Game Theory* 33 (2): 239–254.
- Zhou, Lin. 1990. “On a Conjecture by Gale About One-Sided Matching Problems”. *Journal of Economic Theory* 52 (1): 123–135.

APPENDIX

Rank Efficiency: Modeling a Common Policymaker Objective

Clayton R. Featherstone

The Wharton School, University of Pennsylvania

A Proofs

In several of the proofs that follow, it will be useful to define the **transition graph from assignment x to assignment y** , denoted $G(x \rightarrow y)$, as follows. Start with a vertex representing each agent and object type. Place an edge from (to) object type o to (from) agent a if that agent must take on (offload) probability shares of o in moving from her x allocation to her y allocation. That is, add the edge (o, a) if $y_{ao} > x_{ao}$, and add the edge (a, o) if $y_{ao} < x_{ao}$. Then, add two auxiliary vertices: a source, s , and a sink, t . Place an edge from the source to object type o if there are more shares of o allocated in y than in x , that is, add the edge (s, o) if $\sum_a x_{ao} < \sum_a y_{ao}$. Similarly, place an edge from object type o to the sink if there are more shares of o allocated in x than in y , that is, add the edge (o, t) if $\sum_a x_{ao} > \sum_a y_{ao}$. Finally, remove all vertices with no incident edges. These removed vertices represent agents whose allocations are the same in x and y (agent vertices), object types whose shares are not reallocated in moving from x to y (object type vertices), and the general property that, for each object type, x and y allocate the same amount of shares (the source and sink).

A few properties of $G(x \rightarrow y)$ are immediate. First, the graph is non-empty so long as x and y are distinct assignments. Second, all vertices, except for the sink and source (if they are present), have at least one incoming edge and one outgoing edge. And finally, since $\sum_a \sum_o x_{ao} = \sum_a \sum_o y_{ao} = |A|$, either both the source and sink are in the graph or neither is in the graph. Now, the propositions from the main text will be restated and proved.

A.1 Proofs from [Section 3](#) (Efficiency Concepts)

Proposition 1 (Bogomolnaia and Moulin 2001, Parts a and b).

- (a) *For all ordinal markets, an ordinally efficient assignment must also be ex post efficient.*
- (b) *For some ordinal markets, there are ex post-efficient assignments that fail to be ordinally efficient.*
- (c) *For all ordinal markets, a deterministic assignment is ex post efficient if and only if it is ordinally efficient.*

Proof. For statement (a), by way of contradiction, fix an ordinal market, and assume that some assignment, x , is ordinally efficient, but ex post inefficient. Since it is ex post inefficient, it has an implementation that places positive probability on an ex post-inefficient, deterministic assignment, y . Construct a new random matrix where, in every state of the world where y is realized, instead the ex post dominator of y is realized. The expectation of this random matrix forms an assignment that ordinally dominates x , a contradiction.

For statement (b), the example in [Figure 2](#) provides the proof. Assignment $\frac{1}{2}(w + x)$ is ex post efficient but not ordinally efficient, since it is ordinally dominated by assignment $\frac{1}{2}(y + z)$.

For statement (c), first note that the “if” part is subsumed under statement (a). For the “only if” part, by way of contradiction, fix an ordinal market, and assume that some deterministic assignment, x , is ex post efficient, but ordinally dominated by some other assignment, y . Construct the transition graph, $G(x \rightarrow y)$.^{1a} There are two cases.

There is no cycle in $G(x \rightarrow y)$ Since all nodes but the source and sink have both an incoming and an outgoing edge, by finiteness, the source and sink must be in the graph, and there must be a chain starting at the source and ending at the sink. The object type in that chain that connects to the source, together with an ordered list of the edges from agents to object types in the chain is a claim chain.^{2a} And since x is deterministic, its capacity is 1. Hence, executing it yields a new deterministic assignment in which all agents are weakly better off and some are strictly better off, a contradiction.

There is a cycle in $G(x \rightarrow y)$ Clearly, neither the source nor the sink can be part of the cycle. Hence, an ordered list of the edges from agents to object types in that cycle is a trade cycle. And since x is deterministic, its capacity is 1. Hence, executing it yields a new deterministic assignment in which all agents are weakly better off and some are strictly better off, a contradiction.

Thus, I have established the “only if” part of statement (a), concluding the proof. \square

Proposition 2.

- (a) *For all ordinal markets, a rank-efficient assignment must also be ordinally efficient.*
- (b) *For some ordinal markets, there are ordinally efficient assignments that fail to be rank efficient.*
- (c) *Statement (b) continues to hold if the assignment must be deterministic.*

Proof. This proposition follows immediately from [Proposition 3](#) with partitions \mathcal{P} and \mathcal{P}' being the coarsest and finest partitions of a three-agent student set, respectively, since rank efficiency with respect to the coarsest partition is equivalent to (unadorned) rank efficiency, and rank efficiency with respect to the finest partition is equivalent to ordinal efficiency. \square

Proposition 3.

- (a) *For all ordinal markets, for all partitions of the agent set, \mathcal{P} and \mathcal{P}' , where the latter refines the former, a \mathcal{P} -rank-efficient assignment must also be \mathcal{P}' -rank efficient.*
- (b) *Fix an agent set, \mathcal{A} , with three or more agents, and let \mathcal{P} and \mathcal{P}' be partitions of that set where the latter refines the former.*

1a. The transition graph is defined at the very beginning of [Appendix A](#).

2a. Trade cycles and claim chains are introduced in [Section 5.1](#).

Then, for some ordinal markets with agent set \mathcal{A} , there are \mathcal{P}' -rank-efficient assignments that fail to be \mathcal{P} -rank efficient.

(c) Statement (b) continues to hold if the assignment must be deterministic.

Proof. Throughout this proof, I will denote the elements of \mathcal{P} by \mathcal{A}_i and the elements of \mathcal{P}' by \mathcal{A}_{ij} , where $\mathcal{A}_i = \cup_j \mathcal{A}_{ij}$.

To establish statement (a), fix an ordinal market, and by way of contradiction, assume that some assignment, x , is \mathcal{P} -rank efficient, but \mathcal{P}' -rank dominated by some other assignment, y . Then, $N^y(k; \mathcal{A}_{ij}) \geq N^x(k; \mathcal{A}_{ij})$ holds for all (i, j, k) triples, with the strict inequality holding for at least one triple, (i^*, j^*, k^*) . But, the rank distribution of the set of agents \mathcal{A}_i is just the sum of the rank distributions of the sets of agents \mathcal{A}_{ij} ; therefore, $N^y(k; \mathcal{A}_i) \geq N^x(k; \mathcal{A}_i)$ for all (i, k) pairs, with the strict inequality holding for (i^*, k^*) . Hence, assignment x is \mathcal{P} -rank dominated by assignment y , a contradiction.

To establish statement (b), consider an assignment market where $\mathcal{O} = \{o_1, o_2, \dots, o_{|\mathcal{A}|}\}$ and all object types are unit capacity. Further, let the ranking matrix be given by

$$r = \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & \dots \\ a_1 & \boxed{1} & \boxed{3} & \boxed{2} & 4 & 5 & 6 & \dots \\ a_2 & \boxed{1} & \boxed{2} & \boxed{3} & 4 & 5 & 6 & \dots \\ a_3 & \boxed{3} & \boxed{1} & \boxed{2} & 4 & 5 & 6 & \dots \\ a_4 & 2 & 3 & 4 & 1 & 5 & 6 & \dots \\ a_5 & 2 & 3 & 4 & 5 & 1 & 6 & \dots \\ a_6 & 2 & 3 & 4 & 5 & 6 & 1 & \dots \\ \vdots & \ddots \end{matrix},$$

where the agents and objects besides a_1 , a_2 , and a_3 and o_1 , o_2 , and o_3 are added as needed, and the rankings are extended according to the pattern outside of the highlighted 3×3 submatrix. Further, consider the following two assignments:

$$x = \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & \dots \\ a_1 & \boxed{1} & \boxed{0} & \boxed{0} & 0 & 0 & 0 & \dots \\ a_2 & \boxed{0} & \boxed{1} & \boxed{0} & 0 & 0 & 0 & \dots \\ a_3 & \boxed{0} & \boxed{0} & \boxed{1} & 0 & 0 & 0 & \dots \\ a_4 & 0 & 0 & 0 & 1 & 0 & 0 & \dots \\ a_5 & 0 & 0 & 0 & 0 & 1 & 0 & \dots \\ a_6 & 0 & 0 & 0 & 0 & 0 & 1 & \dots \\ \vdots & \ddots \end{matrix} \quad y = \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & \dots \\ a_1 & \boxed{0} & \boxed{0} & \boxed{1} & 0 & 0 & 0 & \dots \\ a_2 & \boxed{1} & \boxed{0} & \boxed{0} & 0 & 0 & 0 & \dots \\ a_3 & \boxed{0} & \boxed{1} & \boxed{0} & 0 & 0 & 0 & \dots \\ a_4 & 0 & 0 & 0 & 1 & 0 & 0 & \dots \\ a_5 & 0 & 0 & 0 & 0 & 1 & 0 & \dots \\ a_6 & 0 & 0 & 0 & 0 & 0 & 1 & \dots \\ \vdots & \ddots \end{matrix}$$

Again, agents and objects besides a_1 , a_2 , and a_3 and o_1 , o_2 , and o_3 are added as needed, and the assignments are extended according to the pattern outside of the highlighted 3×3 submatrices.

First consider the case where $|\mathcal{P}'| = 2$. Without loss of generality, let agent a_1 be in partition element \mathcal{A}_{11} , and let agents a_2 and a_3 be in \mathcal{A}_{12} . The rest of the agents can be placed into those two partition elements arbitrarily.

Assignment x is \mathcal{P}' -rank efficient. To see this, first note that agent a_1 and agents a_4 through $a_{|\mathcal{A}|}$ must receive their first-choice objects in any assignment that would rank-dominate x . This only leaves switching the assignments of agents a_2 and a_3 , which would not stochastically improve the rank distribution in \mathcal{A}_{12} . But, for \mathcal{P}' to be a refinement of \mathcal{P} , it must be that $\mathcal{A}_1 = \mathcal{A}$. Clearly, assignment x is \mathcal{A} -rank dominated by assignment y , which provides the proof of statement (b), as well as the proof for statement (c), since assignment y is deterministic.

Now, consider the case where $|\mathcal{P}'| \geq 3$. Without loss of generality, let agent a_1 be in partition element \mathcal{A}_{11} , agent a_2 be in \mathcal{A}_{12} , and agent a_3 be in some $\mathcal{A}_{\beta\gamma}$ that isn't \mathcal{A}_{11} or \mathcal{A}_{12} . The rest of the agents can be numbered arbitrarily.

First, I establish that assignment x is \mathcal{P}' -rank efficient. Since all agents get a second choice or better in x , the number of agents who get a second choice is sufficient for the rank distribution. Now, by way of contradiction, assume that some assignment, z , is a \mathcal{P}' -rank dominator of x . Note that all elements of \mathcal{P}' except for \mathcal{A}_{12} and $\mathcal{A}_{\beta\gamma}$ consist solely of agents who get a first choice. Hence, it is only possible to improve the rank distribution of either $\mathcal{A}_{\beta\gamma}$ or \mathcal{A}_{12} , or both.

To improve the rank distribution of $\mathcal{A}_{\beta\gamma}$, agent a_3 must get fewer shares in object o_3 and more shares in object o_2 , which requires that agent a_2 in \mathcal{A}_{12} gets fewer shares in object o_2 and more shares of object o_1 , which requires taking those shares from agent a_1 , a contradiction. Trying to improve the rank distribution of \mathcal{A}_{12} alone leads to a similar contradiction. So, assignment x is \mathcal{P}' -rank efficient. But, it is also \mathcal{P} -rank dominated by assignment y . To see this, consider two cases.

$\beta = 1$: Since $\mathcal{A}_{\beta\gamma} \subseteq \mathcal{A}_1$, it is true that in both x and y , all elements of \mathcal{P} besides \mathcal{A}_1 consist solely of agents who get a first choice. For \mathcal{A}_1 , there are 2 second choice allocations in x and only 1 in y .

$\beta \neq 1$: For both assignments x and y , all elements of \mathcal{P} besides \mathcal{A}_1 and \mathcal{A}_β consist solely of agents who get a first choice. For \mathcal{P}_1 , both x and y admit 1 second choice. For \mathcal{P}_β , assignment x admits 1 second choice, while y admits 0 second choices.

Hence, we have found the desired counterexample. And, since x is deterministic, statement (c) has already been shown. \square

A.2 Proofs from [Section 4](#) (Mechanisms)

Before proving [Proposition 4](#) about random serial dictatorship, it is useful to establish a lemma about deterministic serial dictatorship.

Lemma A.1. Fix an ordinal market, \mathcal{M} . A deterministic assignment, x , is ex post efficient if and only if there exists a deterministic ordering of agents, π , such that $x = SD_\pi(\mathcal{M})$.

Proof. To prove the “if” part, let $x = SD_\pi(\mathcal{M})$. By way of contradiction, assume there is some deterministic assignment, y , that all agents weakly prefer to x and that some set of agents strictly prefer. Construct the transition graph, $G(x \rightarrow y)$, and consider two cases.^{3a}

There is no cycle in $G(x \rightarrow y)$: By finiteness, the source and sink must be in the graph. Follow an edge from the source to an object type, o , and then follow an edge from o to an agent, a . Agent a prefers object type o to his allocation, and what’s more, he could have claimed o when it was his turn in the serial dictatorship, since it wasn’t fully allocated in assignment x , a contradiction.

There is a cycle in $G(x \rightarrow y)$: This cycle can be written as a trade cycle, $\tau = ((a_i, o_i))_{i=1}^n$.^{4a} For each i , it must be that $\pi^{-1}(a_i) > \pi^{-1}(a_{i-1})$; otherwise, a_i would have picked $y_{a_{i-1}}$ (which, by construction, she prefers to x_{a_i}) when it was her turn in the serial dictatorship. But, $\pi^{-1}(a_i) > \pi^{-1}(a_{i-1})$ cannot hold for all edges in the cycle, yielding a contradiction.

So, I have established the “if” part.

To prove the “only if” part, assume x is an ex-post-efficient, deterministic assignment. Denote agent a ’s favorite allocation from set \mathcal{Z} by $c_a(\mathcal{Z})$. Initialize $\mathcal{O}^1 = \mathcal{O}$, $\mathcal{A}^1 = \mathcal{A}$, and $q^1 = q$. Then, consider the recursion

$$\begin{aligned} \mathcal{S}^k &= \{a \in \mathcal{A}^k : x_{a c_a(\mathcal{O}^k)} = 1\}, \\ q_o^{k+1} &= q_o^k - \left| \{a \in \mathcal{S}^k : x_{a o} = 1\} \right|, \text{ for all } o, \\ \mathcal{O}^{k+1} &= \{o \in \mathcal{O}^k : q_o^{k+1} > 0\}, \\ \mathcal{A}^{k+1} &= \mathcal{A}^k \setminus \mathcal{S}^k, \end{aligned}$$

which is understood to terminate when \mathcal{A}^k is empty. Intuitively, \mathcal{S}^1 is the set of agents who got their first choice, \mathcal{S}^2 is the set of agents who got their first choice of what remains after the object types assigned to those in \mathcal{S}^1 are excluded, and so on.

This recursion must terminate. If it didn’t, then there would be some step, k , where no remaining agent gets her first choice of the object types that remain, that is, where \mathcal{S}^k is empty while \mathcal{A}^k is not. But this contradicts the assumption that x was ex post efficient in the first place. To see why, consider a directed graph whose nodes are agents in \mathcal{A}^k , with edges from agent a to agent a' when $x_{a' c_a(\mathcal{O}^k)} = 1$, that is, when agent a' is assigned the object type that agent a ranks first among those in \mathcal{O}^k . Since all nodes have

3a. The transition graph is defined at the very beginning of [Appendix A](#).
4a. Trade cycles and claim chains are introduced in [Section 5.1](#).

an outbound edge, there must be a cycle. Giving every agent in that cycle the object type they point to makes all of them better off, a contradiction.

Given the termination of the recursion, it is clear that serial dictatorship will arrive at x for any agent ordering that places agents in S^1 before agents in S^2 , and so on. \square

Proposition 4 (Svensson 1994). *Fix an ordinal market, \mathcal{M} . An assignment, x , is ex post efficient if and only if there exists some random ordering of agents, Π , such that $x = RSD_{\Pi}(\mathcal{M})$.*

Proof. To prove the “if” part, note that the random matrix $x = SD_{\Pi}(\mathcal{M})$ is an implementation of $RSD_{\Pi}(\mathcal{M})$, and by Lemma A.1, each of its realizations is an ex post-efficient deterministic assignment. Then, by definition, $RSD_{\Pi}(\mathcal{M})$ is an ex post-efficient assignment.

To prove the “only if” part, note that by Lemma A.1, there exists a map from ex post-efficient deterministic assignments to deterministic orderings of agents, φ , that satisfies $z = SD_{\varphi(z)}(\mathcal{M})$ for all ex post-efficient deterministic assignments, z .

Now, take an ex post-efficient assignment, x . It must have an implementation whose support is entirely ex post efficient; call this random variable X . By definition, the random variable $SD_{\varphi(X)}(\mathcal{M})$ is the same as X , and hence their expectations are also equal, that is, $RSD_{\varphi(X)}(\mathcal{M}) = \mathbb{E}[X]$. But, since X is an implementation of x , its expectation is x ; hence, $RSD_{\varphi(X)}(\mathcal{M}) = x$. Therefore, $\varphi(X)$ is the random ordering of agents that establishes the result. \square

Proposition 5 (McLennan 2002; Manea 2008). *Fix an ordinal market, \mathcal{M} . An assignment, x , is ordinally efficient if and only if there exists some valuation matrix, ω , such that $x \in WM_{\omega}(\mathcal{M})$.*

Proof. This follows immediately from Proposition 6 with the partition \mathcal{P} being the finest partition of the student set, since with this choice of partition, \mathcal{P} -rank efficiency is equivalent to ordinal efficiency. \square

For the “only if” part of the proof of Proposition 6, I will need a convex analysis lemma. Recall that the **convex hull** of a set is the intersection of all convex sets that contain it and that a **polytope** is the convex hull of a finite set of points.

Lemma A.2. *Fix a polytope in \mathbb{R}^d , and let P be the set of points in \mathbb{R}^d whose components are all weakly smaller than those of some point in the polytope.*

Then, for each point x in P , exactly one of the following alternatives must hold.

1. *There exists a point μ in \mathbb{R}^d whose components are all strictly positive, such that $\mu \cdot x \geq \mu \cdot z$ for all points z in P .*
2. *There exists a point in P , distinct from x , whose components are all weakly larger than those of x .*

Proof. Both alternatives cannot hold, since the point from the second alternative, whose components are all weakly larger than those of x , has a scalar product with μ from the first alternative that is strictly greater than $\mu \cdot x$, a contradiction.

So, it only remains to show that one of the alternatives must hold. Before doing this, I establish two base results. First, a well-known theorem in convex analysis (Rockafellar 1970, Theorem 19.1) states that P can be represented as the set of solution points, z , to some finite set of inequalities, indexed by i in $\{1, \dots, N\}$, of the form $a^i \cdot z \leq b^i$. Second, for each of these inequalities, a^i must have weakly positive components. To see why, let e_k represent a unit vector along the k th axis of \mathbb{R}^d , and by way of contradiction, assume that, for some i , the k th component of a^i is strictly negative, that is, $a^i \cdot e_k < 0$. Since x is in P , so is $x - \lambda e_k$ for any positive λ . Choose $\lambda = 1 + (a^i \cdot x - b^i)/(a^i \cdot e_k)$, a scalar weakly greater than 1. Then, $x - \lambda e_k$, as a point in P , must obey $a^i \cdot (x - \lambda e_k) \leq b^i$. Rearranging, this yields $a^i \cdot e_k \geq 0$, a contradiction.

Now, I return to showing that one of the two alternatives must always hold. Let \mathcal{I} be the set of indices for inequalities that bind at x , that is, i is in \mathcal{I} if and only if $a^i \cdot x = b^i$. There are two cases to consider.

\mathcal{I} is empty. Define λ_k^i to be the largest positive scalar multiple of e_k that can be added to x without violating inequality i . That is, λ_k^i equals $\sup\{\lambda \geq 0 : a^i \cdot (x + \lambda e_k) \leq b^i\}$. Pick a k where $\inf_i \lambda_k^i$ is finite, and set λ equal to that value. The point $x + \lambda e_k$ is in P , so the second alternative holds.

\mathcal{I} is nonempty. Begin by assuming that the first alternative fails to hold. It must be that for some k , for each i in \mathcal{I} , the k th component of a^i is zero-valued. Otherwise, the first alternative is satisfied by setting $\mu = \sum_{i \in \mathcal{I}} a^i$, a contradiction.

Now, note that the elements of \mathcal{I} all continue to bind at the point $x + \lambda e_k$ for any scalar λ . In addition, any inequalities not indexed by an element of \mathcal{I} , for which $a^i \cdot e_k = 0$ continue to be slack at the point $x + \lambda e_k$ for any scalar λ . Let \mathcal{I}' be the set of indices for which $a^i \cdot e_k > 0$. If \mathcal{I}' is nonempty, set $\lambda = \inf_{i \in \mathcal{I}'} \lambda_k^i$ (where λ_k^i is as defined above); otherwise, set it to an arbitrary positive number. With this choice, all of the defining inequalities of P are either binding or slack at the point $x + \lambda e_k$, which is distinct from x . So, the second alternative holds.

Hence, exactly one of the alternatives holds, as I set out to prove. \square

Proposition 6. Fix an ordinal market, \mathcal{M} , and a partition of the agents in that market, \mathcal{P} . An assignment, x , is \mathcal{P} -rank efficient if and only if there exists some \mathcal{P} -respecting valuation matrix, ω , such that $x \in WM_\omega(\mathcal{M})$.

Proof. Denote the elements of partition \mathcal{P} as $\{A_i\}_{i=1}^{|\mathcal{P}|}$. For the “if” part, note that the ω -welfare objective, evaluated at assignment z , can be rewritten in

terms of the partitioned rank distributions of z according to [Equation 2](#) in the main text.

Now, let assignment x be in $WM_\omega(\mathcal{M})$, and by way of contradiction, take an assignment, y , that \mathcal{P} -rank-dominates x . That is, $N^y(k; \mathcal{A}_i) \geq N^x(k; \mathcal{A}_i)$ for all ranks k and elements \mathcal{A}_i of \mathcal{P} , with the strict inequality holding for some partition element–rank pair. That pair cannot involve the rank $|\mathcal{O}|$, since all agents get their $|\mathcal{O}|$ th choice or better in all assignments. Since the square-bracketed terms in [Equation 2](#) in the main text are strictly positive (by the definition of a valuation), the ω -welfare objective must be larger at y than at x , a contradiction.

For the “only if” part, let \mathcal{X} represent the set of assignments, and define the \mathcal{P} -partitioned rank distribution of assignment z , denoted by $N^z(\mathcal{P})$, to be the $|\mathcal{P}| \times (|\mathcal{O}| - 1)$ -dimensional vector

$$(N^z(1; \mathcal{A}_1), \dots, N^z(|\mathcal{O}| - 1; \mathcal{A}_1), \dots, N^z(1; \mathcal{A}_{|\mathcal{P}|}), \dots, N^z(|\mathcal{O}| - 1; \mathcal{A}_{|\mathcal{P}|})).$$

Note that there are no elements of the form $N^z(|\mathcal{O}|; \mathcal{A}_i)$; these are omitted because they are equal to $|\mathcal{A}_i|$, regardless of the assignment z . Similar to the notation introduced for \mathcal{P} -respecting valuation matrices in [Section 4.2](#), when I reference the (\mathcal{A}_i, k) element of a vector in $\mathbb{R}^{|\mathcal{P}| \times (|\mathcal{O}| - 1)}$, I mean its $[(i - 1)(|\mathcal{O}| - 1) + k]$ th component.

Now, denote by $\mathcal{N}(\mathcal{P})$ the set of all \mathcal{P} -partitioned rank distributions that can be generated by some assignment. By the Budish et al. (2013) result mentioned in [Section 2](#), the set of assignment is the convex hull of the finite set of deterministic assignments. And by the linearity of the definition of the \mathcal{P} -partitioned rank distribution, the \mathcal{P} -partitioned rank distribution of the convex combination of two deterministic assignments is equal to the convex combination of those assignments’ \mathcal{P} -partitioned rank distributions. That is, if y and z are deterministic assignments and $\lambda \in [0, 1]$, then $N^{\lambda y + (1 - \lambda)z}(\mathcal{P}) = \lambda N^y(\mathcal{P}) + (1 - \lambda)N^z(\mathcal{P})$. So, the set $\mathcal{N}(\mathcal{P})$ is the convex hull of the set of deterministic assignments’ \mathcal{P} -partitioned rank distributions. That is, $\mathcal{N}(\mathcal{P})$ is a polytope.

In this context, the second alternative in [Lemma A.2](#) is that an assignment is \mathcal{P} -rank dominated. So, if an assignment is \mathcal{P} -rank efficient, the first alternative of [Lemma A.2](#) must hold, that is, the assignment’s \mathcal{P} -partitioned rank distribution maximizes [Equation 2](#) in the main text if $\omega_{\mathcal{A}_i k} - \omega_{\mathcal{A}_i(k+1)}$ is equal to $\mu_{\mathcal{A}_i k}$. This relation is the reason for the emphasis on strictly positive μ in [Lemma A.2](#); if μ had a zero-valued component, then ω wouldn’t be a valid valuation matrix.

Hence, for each element of the partition, $\omega_{\mathcal{A}_i}$, the valuation for agents in \mathcal{A}_i , is given componentwise according to $\omega_{\mathcal{A}_i k} = \sum_k^{|\mathcal{O}| - 1} \mu_{\mathcal{A}_i j}$ for k in $\{1, \dots, |\mathcal{O}| - 1\}$ (recall that $\omega_{\mathcal{A}_i |\mathcal{O}|} = 0$, by definition). By construction, this valuation matrix, ω , is \mathcal{P} -rank respecting and $x \in WM_\omega(\mathcal{M})$. \square

Proposition 7. *Fix an ordinal market, \mathcal{M} . An assignment, x , is rank efficient if and only if there exists some valuation vector, v , such that $x \in RV_v(\mathcal{M})$.*

Proof. This follows immediately from [Proposition 6](#) with the partition \mathcal{P} being the coarsest partition of the student set, since with this choice of partition, \mathcal{P} -rank efficiency is equivalent to rank efficiency. \square

A.3 Proofs from [Section 5](#) (A Model of Tinkering)

In moving on to the propositions that concern the tinkering process of [Section 5](#), it is useful to slightly generalize the definition of executing a claim chain or trade cycle introduced in [Section 5.1](#). In the main text, the fractional share passed along a trade cycle or claim chain is always assumed to be the relevant capacity. Of course, it is entirely feasible to pass a fractional share that is less than the capacity, a number which I will refer to as the **volume** of the execution. Obviously, the execution volume cannot be larger than the capacity of the trade cycle or claim. When a specific volume is not mentioned, the volume should be understood to be the capacity, in keeping with the definition in the main text. When a different volume is specified, all definitions of executing a trade cycle or claim chain in [Section 5.1](#) remain the same, except that c takes on the specified volume.

Now, for the “if” part of the proof of [Proposition 8](#), the following lemma will prove useful. It demonstrates that transitioning from one assignment to another can be achieved by executing a sequence of trade cycles and claim chains.

Lemma A.3. *Fix an assignment problem and two assignments, x and z . There exists a finite list of ordered triples $((y^i, \sigma^i, c^i))_{i=1}^n$ such that*

- (a) *The first assignment, y^1 , is equal to x ,*
- (b) *For all i in $\{1, \dots, n\}$, the trade cycle or claim chain in y^i represented by σ^i , when executed at volume c^i , yields y^{i+1} (where y^{n+1} is defined to be z), and*
- (c) *For all i in $\{1, \dots, n\}$, the trade cycle or claim chain represented by σ^i is also a trade cycle or claim chain in x .*

Proof. First, I show that in any assignment, y^i , that is distinct from z , I can find a trade cycle or claim chain that, when executed at an appropriately chosen volume, yields a new assignment, y^{i+1} , such that the transition graph $G(y^{i+1} \rightarrow z)$ has fewer edges than $G(y^i \rightarrow z)$.^{5a} There are two cases to consider.

There is a cycle in $G(y^i \rightarrow z)$. An ordered list of the edges in that cycle that are outgoing from an agent corresponds to a trade cycle, $((a_j, o_j))_{j=1}^m$. Choose a volume equal to

$$\min \left\{ \min_{j \in \{1, \dots, m\}} \{y_{a_j o_j}^i - z_{a_j o_j}\}, \min_{j \in \{1, \dots, m\}} \{z_{a_j o_{j-1}} - y_{a_j o_{j-1}}^i\} \right\},$$

where o_0 is defined to be o_m . Let y^{i+1} be the assignment yielded by executing the trade cycle at that volume in y^i . The volume is designed

5a. The transition graph is defined at the very beginning of [Appendix A](#).

to be the largest that won't cause any agent to overshoot their allocation in z , that is, executing the trade cycle will leave $G(y^{i+1} \rightarrow z)$ with a strict subset of the set of edges of $G(y^i \rightarrow z)$. Strictness follows from the fact that the min will bind on at least one of the edges in $G(y^i \rightarrow z)$. The quantities in the first interior min represent the amount of shares of object type o_j that agent a_j needs to offload to get to her allocation of o_j in assignment z , while the quantities in the second interior min represent the number of shares of object type o_{j-1} that agent a_j needs to take on to get to her allocation of o_{j-1} in z . By construction, $G(y^{i+1} \rightarrow z)$ has at least one less edge than $G(y^i \rightarrow z)$.

There is no cycle in $G(y^i \rightarrow z)$. Since y^i is assumed distinct from z , there must be a chain that starts at the source and ends at the sink. The object type connected to the source, o_s , and an ordered list of the agent to object type edges in the chain is a claim chain in y^i , which I denote by $(o_s, ((a_j, o_j))_{j=1}^m)$. Choose a volume equal to

$$\min \left\{ \min_{j \in \{1, \dots, m\}} \{y_{a_j o_j}^i - z_{a_j o_j}\}, \min_{j \in \{1, \dots, m\}} \{z_{a_j o_{j-1}} - y_{a_j o_{j-1}}^i\}, \sum_a y_{a o_s}^i - \sum_a z_{a o_s}, \sum_a z_{a o_m} - \sum_a y_{a o_m}^i \right\},$$

where o_0 is defined to be o_s . Let y^{i+1} be the assignment yielded by executing the claim chain at that volume in y^i . Again, the volume is designed to be the largest that won't cause any agent to overshoot their allocation in z , that is, executing the trade cycle will leave $G(y^{i+1} \rightarrow z)$ with a strict subset of the set of edges of $G(y^i \rightarrow z)$. The first line of the expression has the same interpretation as in the "cycle" case above, while the second is the same restriction applied to the source and the sink. Again, by construction, the transition graph $G(y^{i+1} \rightarrow z)$ has at least one less edge than $G(y^i \rightarrow z)$.

Now, initialize $y^1 = x$, and choose σ^i and c^i to generate y^{i+1} . If y^i is distinct from z , then it is clear that $G(y^{i+1} \rightarrow z)$ has strictly fewer edges than $G(y^i \rightarrow z)$. By finiteness, this process must terminate with z .

Finally, note that since the set of edges in $G(y^i \rightarrow z)$ is decreasing in i , any trade cycle or claim chain in y^i is also a trade cycle or claim chain in x . \square

Now that the groundwork has been laid, I can prove [Proposition 8](#).

Proposition 8. *Fix an ordinal market. An assignment, x , is rank efficient if and only if there exists a valuation vector, v , such that there is no trade cycle or claim chain in x that is v -welfare improving.*

Proof. Let the fixed ordinal market be \mathcal{M} . I begin by proving the "only if" part. If x is rank efficient, then by [Proposition 7](#), there must exist a valuation vector, v , such that $x \in RV_v(\mathcal{M})$. I claim that there are no v -welfare-improving trade cycles or claim chains in x . By way of contradiction, assume

otherwise. Let σ denote the trade cycle or claim chain in x that is v -rank improving. If the capacity of σ is κ , then the v -welfare objective, evaluated at x , is less than that same function evaluated at the assignment that results from implementing σ in x : the difference is $\kappa \Delta(\sigma, v)$. So, $x \notin RV_v(\mathcal{M})$, a contradiction.

Now, I prove the “if” part. By way of contradiction, assume that, for all valuations v , there are no v -welfare-improving trade cycles or claim chains in x , but that x is not rank efficient. Fix any valuation, v . By [Proposition 7](#), it must be that $x \notin RV_v(\mathcal{M})$. Pick some $z \in RV_v(\mathcal{M})$. By [Lemma A.3](#), there is a finite list of ordered triples, $((y^i, \sigma^i, c^i))_{i=1}^n$ where y^1 equals x such that, for all i in $\{1, \dots, n\}$, the trade cycle or claim chain in y^i represented by σ^i , when executed at volume c^i , yields y^{i+1} (where y^{n+1} is defined to be z).

Now, consider the v -welfare objective. This function increases by $c_i \Delta(\sigma_i; v)$ when it moves from y^i to y^{i+1} . So, the function evaluated at z is $\sum_{i=1}^n c_i \Delta(\sigma_i; v)$ greater than the function evaluated at x . Since z is in $RV_v(\mathcal{M})$, while x is not, it must be that $\sum_{i=1}^n c_i \Delta(\sigma_i; v) > 0$. But then there must exist some i^* such that $\Delta(\sigma_{i^*}; v) > 0$. But then by part (c) of [Lemma A.3](#), the trade cycle or claim chain represented by σ_{i^*} is in x . And, it is v -rank improving, a contradiction. \square

The example in [Figure A.1](#) establishes that absence of rank-improving cycles is not sufficient for rank efficiency. Formally,

Proposition A.1. *For all ordinal markets, an assignment that is rank efficient has no rank-improving trade cycles or claim chains. However, there exist ordinal markets such that an assignment with no rank-improving trade cycles or claim chains can fail to be rank efficient.*

Proof. The first statement is an immediate consequence of [Proposition 8](#), since a rank-improving trade cycle or claim chain is also v -welfare improving for all valuations, v . For the second statement, consider the seven-agent–seven-object type example in [Figure A.1](#), where each object type has unit capacity. I will only consider assignments where all agents get a 3rd choice or better. Given the rankings in part (a), this requires that agents a_1 through a_4 receive objects o_1 through o_4 and agents a_5 through a_7 receive objects o_5 through o_7 .

In part (b), assignment x is rank dominated by assignment y . Clearly, there are no rank-improving claim chains, since all objects are already claimed. There are also no rank-improving trade cycles. To see this, first note that any such trade cycle could only involve agents from a subset of $\{a_1, a_2, a_3, a_4\}$ or a subset of $\{a_5, a_6, a_7\}$. A rank-improving trade cycle involving agents from the former subset would improve the rank distribution of the sub-ordinal market involving those agents and objects o_1 through o_4 . Similarly, a rank-improving trade cycle involving agents from the latter subset would improve the rank distribution of the sub-ordinal market

$$r \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & o_7 \\ \begin{matrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ a_7 \end{matrix} & \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 1 & 2 & 4 & 3 & 5 & 6 & 7 \\ 3 & 1 & 2 & 4 & 5 & 6 & 7 \\ 3 & 4 & 1 & 2 & 5 & 6 & 7 \\ 4 & 5 & 6 & 7 & 1 & 2 & 3 \\ 4 & 5 & 6 & 7 & 3 & 1 & 2 \\ 4 & 5 & 6 & 7 & 1 & 2 & 3 \end{pmatrix} \end{matrix}$$

(a) Rankings

$$x \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & o_7 \\ \begin{matrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ a_7 \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}$$

$$y \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 & o_5 & o_6 & o_7 \\ \begin{matrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ a_6 \\ a_7 \end{matrix} & \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \end{matrix}$$

$$N^x = \begin{pmatrix} \text{1st} & \text{2nd or} & \text{3rd or} \\ \text{choice} & \text{better} & \text{better} \\ 3 & 6 & 7 \end{pmatrix}$$

$$N^y = \begin{pmatrix} \text{1st} & \text{2nd or} & \text{3rd or} \\ \text{choice} & \text{better} & \text{better} \\ 4 & 6 & 7 \end{pmatrix}$$

(b) Two assignments and their rank distributions: y rank dominates x

$$\begin{aligned} \tau_1 &= \{(a_2, o_2), (a_3, o_3), (a_4, o_4)\} & \tau_2 &= \{(a_6, o_6), (a_7, o_7)\} \\ \Delta(\tau_1; v) &= 2v_1 - 3v_2 + v_3 & \Delta(\tau_2; v) &= -v_1 + 2v_2 - v_3 \end{aligned}$$

(c) The two disjoint trade cycles required to move from x to y . Neither is rank improving.

Figure A.1: Example for proof of [Proposition A.1](#)

NOTES: The assignment problem described in this figure has agents $\mathcal{A} = \{a_1, a_2, a_3, a_4, a_5, a_6, a_7\}$, preferences as listed in part (a), and object types $\mathcal{O} = \{o_1, o_2, o_3, o_4, o_5, o_6, o_7\}$, each with unit capacity.

involving those agents and objects o_5 through o_7 . I will argue that neither scenario is possible.

First, consider the sub-ordinal market involving agents a_1 through a_4 and objects o_1 through o_4 . In assignment x , those agents get one 1st choice and three 2nd choices. So, to improve the rank distribution, at the very least, all agents must get a 2nd choice or better. Clearly, a_1 and a_2 must get o_1 and o_2 ; otherwise, at least one of them would be relegated to a 3rd choice. But then, a_3 must take o_3 to avoid getting a 3rd choice, which leaves o_4 for a_4 , her 2nd choice. That is, if agents a_1 through a_4 get a 2nd or better choice, then at most one of them can get a 1st choice. Hence, the sub-assignment in x involving agents a_1 through a_4 is rank efficient for the sub-ordinal market we are considering. Then, by the first statement in this proposition, there can't be a rank-improving cycle involving agents a_1 through a_4 .

Now, consider the sub-ordinal market involving agents a_5 through a_7 and objects o_5 through o_7 . In assignment x , those agents get two 1st choices and a 3rd choice. Clearly, it is impossible to give more than two 1st choices. So, to improve the rank distribution, we need all agents to get a 2nd choice or better. But this is impossible: without loss of generality, if a_5 gets o_5 , then a_7 has to get o_7 , her 3rd choice. Hence, the sub-assignment in x involving agents a_5 through a_7 is rank efficient for the sub-ordinal market we are considering. Then, by the first statement in this proposition, there can't be a rank-improving cycle involving agents a_5 through a_7 .

Hence, there are no rank-improving trade cycles or claim chains in assignment x . And yet, as I established initially, x is rank dominated by y , establishing the second statement of the proposition. \square

Incidentally, part (c) of [Figure A.1](#) lists the two disjoint trade cycles, τ_1 and τ_2 , required to move from assignment x to assignment y . It follows from [Proposition 8](#) that for any rank vector v , either $\Delta(\tau_1; v)$ or $\Delta(\tau_2; v)$ must be strictly positive. This is true, for if both were weakly negative, we would have $v_2 \in [\frac{2}{3}v_1 + \frac{1}{3}v_3, \frac{1}{2}v_1 + \frac{1}{2}v_3]$, which is impossible when $v_1 > v_3$, as it is in a valuation vector.

Proposition 9. *Fix an ordinal market, \mathcal{M} . The v -Sequential Improvements Algorithm ([Algorithm 1](#)) terminates in a finite number of steps and outputs a rank-efficient assignment that is a member of $RV_v(\mathcal{M})$.*

Proof. Begin by taking the minimum of $\Delta(\sigma; v)$ over all possible trade cycles and claim chains in all assignments, σ , such that $\Delta(\sigma; v) > 0$. By finiteness, this minimum is attained by some σ and is equal to some value, $\underline{\Delta}$, that is strictly positive.

Now, note that if x^t is deterministic, then all trade cycles and claim chains in x^t have capacity 1, and hence x^{t+1} is also deterministic. Since x^0 is deterministic by assumption, an induction argument establishes that if the algorithm doesn't terminate by outputting x^t , then σ^t has a capacity of 1 and x^{t+1} is deterministic.

Then, by way of contradiction, assume that the algorithm fails to terminate for some inputs x^0 and v . In moving from x^t to x^{t+1} , the v -welfare objective increases by at least $\underline{\Delta}$. Thus the v -welfare objective of x^t increases without bound in t . This contradicts the fact that $RV_v(\mathcal{M})$ is nonempty, which is clear from its defining LP (Equation 1).

Denote the terminal assignment by x^T . By way of contradiction, assume that x^T is not in $RV_v(\mathcal{M})$. By Lemma A.3 and the definition of the v -rank-value mechanism, there exists a v -welfare improving trade cycle or claim chain in x^T , which contradicts the assumption that it is terminal. And since x^T is in $RV_v(\mathcal{M})$, by Proposition 7, it must also be rank efficient. \square

A.4 Proofs from Section 6 (Truthful Revelation)

Proposition 10. *No ordinal assignment mechanism is both rank efficient and strategyproof.*

Proof. Consider agent a_1 and the market realization she faces in Figure A.2. If indeed, assignments x and \tilde{x} in part (c) are the respective unique rank-efficient assignments when the submitted rankings are $(r_{a_1}; \tilde{r}_{-a_1})$ and $(\tilde{r}_{a_1}; \tilde{r}_{-a_1})$, then the proof is complete, since by the true preferences of agent a_1 , listed in part (b), she does unambiguously better to report \tilde{r}_{a_1} , which results in assignment \tilde{x} , where she gets object o_1 , which she ranks 1st. Truth-telling results in assignment x , where she gets object o_4 , which she ranks 2nd.

Now I show that x and \tilde{x} are what I claim. I begin by showing that x is the unique rank-efficient assignment when the rankings are $(r_{a_1}; \tilde{r}_{-a_1})$. Given x and $(r_{a_1}; \tilde{r}_{-a_1})$, there are three 1st-choice allocations and one 2nd-choice allocation. Clearly, no assignment can give four 1st choices, since both agents a_1 and a_2 share a top-ranked object, namely o_1 . So, the best an alternative assignment could do is match the rank distribution of x . To do this necessarily entails having agents a_1 and a_2 swap, but this introduces a 3rd-choice allocation. So x is indeed the unique rank-efficient assignment when the rankings are $(r_{a_1}; \tilde{r}_{-a_1})$.

Now, I show that \tilde{x} is the unique rank-efficient assignment when the rankings are $(\tilde{r}_{a_1}; \tilde{r}_{-a_1})$. Assignment \tilde{x} allocates three 1st choices and one 3rd choice. Again, no assignment can give four 1st choices, since both agents a_1 and a_2 share o_1 as a top-ranked object. It is also impossible for more than three agents to get a 2nd-choice or better allocation, since no agent ranks object o_4 higher than 3rd. So, the best an alternative assignment could do is match the rank distribution of \tilde{x} . To do this necessarily entails having agents a_1 and a_2 swap, but this introduces a 4th-choice allocation. So, \tilde{x} is indeed the unique rank-efficient assignment when the ranking matrix is $(\tilde{r}_{a_1}; \tilde{r}_{-a_1})$, and the proof is complete. \square

Now, I will prove the various claims made in Section 6.2.1 of the main text, which introduced the tiebreaker procedure. To do so, I will use a new

$$\tilde{r}_{-a_1} \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 \\ a_2 & \left(\begin{array}{cccc} 1 & 2 & 4 & 3 \\ 4 & 1 & 2 & 3 \\ 2 & 3 & 1 & 4 \end{array} \right) \\ a_3 & \\ a_4 & \end{matrix}$$

(a) Reported rankings of agents a_2 , a_3 , and a_4

$$r_{a_1} \equiv (1 \ 3 \ 4 \ 2) \quad \tilde{r}_{a_1} \equiv (1 \ 2 \ 3 \ 4)$$

(b) For agent a_1 , the true ranking is r_{a_1} ,
while the alternative report is \tilde{r}_{a_1} .

$$x \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 \\ a_1 & \left(\begin{array}{cccc} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right) \\ a_2 & \\ a_3 & \\ a_4 & \end{matrix} \quad \tilde{x} \equiv \begin{matrix} & o_1 & o_2 & o_3 & o_4 \\ a_1 & \left(\begin{array}{cccc} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right) \\ a_2 & \\ a_3 & \\ a_4 & \end{matrix}$$

(c) Assignment x is the unique rank-efficient assignment
given rankings $(r_{a_1}; \tilde{r}_{-a_1})$; assignment \tilde{x} is the unique
rank-efficient assignment given rankings $(\tilde{r}_{a_1}; \tilde{r}_{-a_1})$.

Figure A.2: Market realization for the proof of [Proposition 10](#)

NOTES: This figure represents a market realization from the perspective of agent a_1 . The object types are $\mathcal{O} = \{o_1, o_2, o_3, o_4\}$, each with unit capacity, and the true ranking of agent a_1 is r_{a_1} , as listed in part (b). The other agents are $\mathcal{A}_{-a_1} = \{a_2, a_3, a_4\}$, and their reported rankings are \tilde{r}_{-a_1} , as listed in part (a).

notation: let $W_\omega(x)$ denote the value of the ω -welfare objective evaluated at assignment x .

Proposition A.2. Fix an ordinal market, \mathcal{M} , and a valuation matrix, ω .

- (a) The set of assignments $WM_\omega(\mathcal{M})$ is the convex hull of a finite set of deterministic assignments.
- (b) For any deterministic ordering of agents, the output of the tiebreaker algorithm ([Algorithm 2](#)) contains exactly one deterministic assignment.
- (c) For any assignment in $WM_\omega(\mathcal{M})$, there exists a random ordering such that the corresponding tiebreaker yields that assignment.

Proof. The proof of part (a) follows from the showing that if x is in $WM_\omega(\mathcal{M})$, and X is an implementation of x , then all deterministic assignments in the support of X must also be in $WM_\omega(\mathcal{M})$.

Let the implementation X select deterministic assignment y^i with positive probability λ^i for i in $\{1, \dots, n\}$. There's nothing to show if $n = 1$, so assume $n \geq 2$. By linearity, $W_\omega(x) = \lambda^i \sum_i W_\omega(y^i)$. Let V^* equal $W_\omega(x)$.

Now, by way of contradiction, and without loss of generality, assume that y^1 is has an ω -welfare objective that is strictly less than V^* . Then, define $\Lambda = \sum_{j>1} \lambda^j$ and $z = \sum_{i>1} (\lambda^i / \Lambda) y^i$. It must be that $V^* = \Lambda W_\omega(z) + \lambda^1 W_\omega(y^1)$, which means that $W_\omega(z) > V^*$. This contradicts that x was in $WM_\omega(\mathcal{M})$ in the first place.

The statement in part (b) could only fail if $\mathcal{X}_{|\mathcal{A}|}$ in [Algorithm 2](#) were either empty or contained at least two assignments. The former case cannot be, since the algorithm starts with a nonempty set (the defining LP of the welfare maximization mechanism, [Equation 1](#), maximizes a linear function on a compact set) and cannot ever eliminate all assignments in \mathcal{X}_t at a given step t . To rule out the second possibility, by way of contradiction, assume there are two assignments in $\mathcal{X}_{|\mathcal{A}|}$. Then for some t , agent $\pi^{-1}(t)$ gets a different object type in each, which means that at step t of the algorithm, an assignment remained that failed to give agent $\pi^{-1}(t)$ his most preferred object type, a contradiction.

To show part (c), for some valuation matrix, ω , let $\mathcal{X} = WM_\omega(\mathcal{M})$, and pick some $x \in \mathcal{X}$. By [Proposition 7](#), the assignment x is ordinally efficient, which by [Proposition 1](#), means that it is also ex post efficient. Then, by [Proposition 4](#), there is a random ordering of agents, Π , such that $RSD_\Pi(\mathcal{M}) = x$. I will now show that the tiebreaker procedure over $WM_\omega(\mathcal{M})$ given random ordering Π will also generate x , that is, $b_\Pi(\mathcal{X}) = x$.

I will prove this by proving that, for any deterministic tiebreaker, π , in the support of Π , if $y = SD_\pi(\mathcal{M})$, then $y = b_\pi(\mathcal{X})$ as well. To help with notation, note that the tiebreaker algorithm ([Algorithm 2](#)) can be slightly altered to yield serial dictatorship by initializing \mathcal{X}_0 to the set of all deterministic assignments in ordinal market \mathcal{M} (instead of just the set of deterministic assignments in the more general input, \mathcal{X}). I will denote the remaining assignments after step t under serial dictatorship and the tiebreaker by \mathcal{X}_t^{SD} and \mathcal{X}_t^b , respectively. The proof is by induction on rounds of the tiebreaker.

Base step ($y \in \mathcal{X}_0^b$). By the same rationale used in the proof of part (a), if $x \in \mathcal{X}_0^b$, then $y \in \mathcal{X}_0^b$ as well.

Induction step ($y \in \mathcal{X}_t^b \Rightarrow y \in \mathcal{X}_{t+1}^b$). By way of contradiction, assume $y \in \mathcal{X}_t^b$ but that $y \notin \mathcal{X}_{t+1}^b$. By the definition of the tiebreaker, \mathcal{X}_t^b contains the elements of $WM_\omega(\mathcal{M})$ that give agents $\pi(1)$ through $\pi(t)$ the object types they get in y . By the definition of serial dictatorship, \mathcal{X}_t^{SD} contains all deterministic assignments that give agents $\pi(1)$ through $\pi(t)$ the object type they get in y . Hence, $\mathcal{X}_t^b \subseteq \mathcal{X}_t^{SD}$.

This means that if $y \notin \mathcal{X}_{t+1}^b$, then agent $\pi(t+1)$ is getting an object type under the tiebreaker that she prefers to what she gets in assignment y . But this preferred object type is also available to agent $\pi(t+1)$ in the serial dictatorship algorithm. That she fails to get it in y provides the contradiction.

So, by induction, $y = b_\pi(\mathcal{X})$, as I sought to prove. \square

Now, I prove the propositions concerning truth-telling, truncation, and extension in low-information environments. Going forward, I will need to consider the v -welfare objective evaluated with different rankings. Let $W_v(x, r)$ denote the v -welfare objective of assignment x when the ranking profile is r .

Further, I will consider rankings and assignments where I apply the switching operator to just the component corresponding to agent a . Define the assignment $x^{o \leftrightarrow o', a}$ to be the same as assignment x , except that the allocations of object types o and o' for agent a are switched. That is, $x^{o \leftrightarrow o', a}$ is the same as x , except that $x_{ao}^{o \leftrightarrow o', a} = x_{ao'}$ and $x_{ao'}^{o \leftrightarrow o', a} = x_{ao}$. The same notation can also be applied to a ranking matrix.

Now, to prove [Proposition 11](#), I must prove [Proposition A.3](#). A lemma will prove helpful for doing so.

Lemma A.4. *Assume that agent a strictly prefers object type o' to object type o in ranking r_a . If $x \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $y \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow o', a}, q)$, then $x_{ao} - x_{ao'} \leq y_{ao} - y_{ao'}$.*

Proof. Begin by noting that

$$\begin{aligned} W_v(x, \tilde{r}) &= W_v(x, \tilde{r}^{o \leftrightarrow o', a}) - (x_{ao} - x_{ao'}) (v_{\tilde{r}_{ao'}} - v_{\tilde{r}_{ao}}), \text{ and} \\ W_v(y, \tilde{r}) &= W_v(y, \tilde{r}^{o \leftrightarrow o', a}) - (y_{ao} - y_{ao'}) (v_{\tilde{r}_{ao'}} - v_{\tilde{r}_{ao}}). \end{aligned}$$

Since by optimality, the inequality $W_v(x, \tilde{r}) \geq W_v(y, \tilde{r})$ holds, it must be that

$$\begin{aligned} W_v(x, \tilde{r}^{o \leftrightarrow o', a}) &\geq W_v(y, \tilde{r}^{o \leftrightarrow o', a}) \\ &\quad + [(x_{ao} - x_{ao'}) - (y_{ao} - y_{ao'})] (v_{\tilde{r}_{ao'}} - v_{\tilde{r}_{ao}}). \end{aligned}$$

But, since optimality also implies that $W_v(y, \tilde{r}^{o \leftrightarrow o', a}) \geq W_v(x, \tilde{r}^{o \leftrightarrow o', a})$ holds, and $(v_{\tilde{r}_{ao'}} - v_{\tilde{r}_{ao}})$ is positive, it must be that $x_{ao} - x_{ao'} \leq y_{ao} - y_{ao'}$, which is what I set out to prove. \square

Proposition A.3. *Consider an agent, a , who faces a rank-value mechanism and who knows the set of object types is \mathcal{O} . For two object types, o and o' in \mathcal{O} , let her truly prefer object type o' to object type o , and let her beliefs be $\{o, o'\}$ -symmetric.*

If o' is ranked higher than o in the reported ranking \tilde{r}_a , then agent a weakly prefers (in the first-order stochastic sense) her expected allocation from \tilde{r}_a to her expected allocation from $\tilde{r}_a^{o \leftrightarrow o'}$.

Proof. Consider agent a submitting \tilde{r}_a . Let x represent the assignment selected when the market realization is $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$. Similarly, let z represent the assignment selected when the scenario is market realization is $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}^{o \leftrightarrow o'}, q^{o \leftrightarrow o'})$. Given her $\{o, o'\}$ -symmetric beliefs, agent a thinks both of these scenarios are equally likely, meaning that her expectation for the assignment (conditional on one of these two market realizations) is $(x + z)/2$.

Now, consider what happens in these market realizations if agent a submits $\tilde{r}_a^{o \leftrightarrow o'}$ instead of \tilde{r}_a . Respectively, the mechanism must return $z^{o \leftrightarrow o'}$ and $x^{o \leftrightarrow o'}$. Then, given the $\{o, o'\}$ -symmetric beliefs of agent a , she thinks both of these realized revelation problems are equally likely, meaning that her expectation for the assignment (conditional on one of the two market realizations) is $(x^{o \leftrightarrow o'} + z^{o \leftrightarrow o'})/2$.

So, in moving from submitting \tilde{r}_a to submitting $\tilde{r}_a^{o \leftrightarrow o'}$, the assignment (conditional on one of the two market realizations) changes by $\delta = [(x^{o \leftrightarrow o'} - x) + (z^{o \leftrightarrow o'} - z)]/2$. The components of this matrix are zero except for those in the o and o' columns. Therefore, the change to the o allocation of agent a is $\delta_{ao} = [(x_{ao'} - x_{ao}) + (z_{ao'} - z_{ao})]/2$, while the change to her o' allocation is $\delta_{ao'} = -\delta_{ao}$. If $\delta_{ao} \geq 0$, then this is a first-order stochastic loss.

In the interest of aligning with the notation of [Lemma A.4](#), define $y = z^{o \leftrightarrow o'}$. That way, $x \in b_\pi \circ RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $y \in b_\pi \circ RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow o'}, q)$. Then, [Lemma A.4](#) implies that $x_{ao} - x_{ao'} \leq y_{ao} - y_{ao'}$. Rewriting δ_{ao} in terms of y yields $\delta_{ao} = [-(x_{ao} - x_{ao'}) + (y_{ao} - y_{ao'})]/2$. So, $\delta_{ao} \geq 0$ —moving from submitting \tilde{r}_a to submitting $\tilde{r}_a^{o \leftrightarrow o'}$ is indeed a first-order stochastic loss, conditional on one of the two market realizations I have focused on thus far.

The entire space of possible market realizations can be partitioned into pairs related by the $o \leftrightarrow o'$ operation. Hence, for any distribution over those pairs, the logic of the previous paragraph shows that moving from submitting \tilde{r}_a to submitting $\tilde{r}_a^{o \leftrightarrow o'}$ is a first-order stochastic loss. So, agent a prefers her allocation (in the first-order stochastic sense) when she submits \tilde{r}_a to her allocation when she submits $\tilde{r}_a^{o \leftrightarrow o'}$, as I sought to prove. \square

The proof of [Proposition 11](#) follows directly from [Proposition A.3](#).

Proposition 11. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let \mathcal{S} be a subset of \mathcal{O} , and let the agent have \mathcal{S} -symmetric beliefs.*

Then, for any report that isn't truthful about \mathcal{S} , there exists another report that is truthful about \mathcal{S} such that the agent weakly prefers (in the first-order stochastic sense) her expected allocation from the latter report to her expected allocation from the former.

Proof. Pick any initial report that isn't truthful about \mathcal{S} , and call it $\tilde{r}_a(0)$. For any report, $\tilde{r}_a(t)$ that isn't truthful about \mathcal{S} , there must be two object types, o and o' in \mathcal{S} , out of order. Set $\tilde{r}_a(t+1) = \tilde{r}_a(t)^{o \leftrightarrow o'}$. By [Proposition A.3](#), she must weakly prefer (in the first-order stochastic sense) her allocation when she reports $\tilde{r}_a(t+1)$ to her allocation when she reports $\tilde{r}_a(t)$.

This sequence must terminate at some $\tilde{r}_a(T)$, since there are only a finite number of swaps that can be done. And the sequence can only terminate at a report that is truthful about \mathcal{S} . So, $\tilde{r}_a(T)$ is a ranking that is truthful about \mathcal{S} and that gives her an allocation that she weakly prefers (in the first-order stochastic sense) to what she gets when she reports $\tilde{r}_a(0)$, as I set out to show. \square

The proof of [Proposition 12](#) follows directly from [Proposition 11](#).

Proposition 12. Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let her have \mathcal{O} -symmetric beliefs.

Then, the agent weakly prefers (in the first-order stochastic sense) her expected allocation from truth-telling to her expected allocation from any other report.

Proof. The only report that is truthful about \mathcal{O} is the true ranking. Hence, by Proposition 11, truth-telling gives an allocation that is weakly preferred (in the first-order stochastic sense) to the allocation that results from any other report. \square

To prove Proposition 13, I will use Proposition A.4, which I prove below. That proof requires another lemma.

Lemma A.5. Consider an agent, a , who knows the object type set is \mathcal{O} and that it contains \emptyset , and fix an \mathcal{O} -conformable market realization, $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$. Let o be the object type directly below \emptyset in some ranking, \tilde{r}_a , that is, $\tilde{r}_{ao} = \tilde{r}_{a\emptyset} + 1$.

If there exists some \hat{x} in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ such that $\hat{x}_{a\emptyset} = 0$ and some \hat{y} in $RV(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$ such that $\hat{y}_{a\emptyset} = 0$, then the following statements hold.

- (a) If y is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$ and $y_{a\emptyset} = 0$, then $y^{o \leftarrow \emptyset, a}$ is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$.
- (b) If x is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $x_{a\emptyset} = 0$, then x is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$.

Proof. To show part (a), by way of contradiction, assume that y is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$ and $y_{a\emptyset} = 0$, but that $y^{o \leftarrow \emptyset, a}$ is not in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$. Since y is feasible and $y_{a\emptyset} = 0$, it must also be that $y^{o \leftarrow \emptyset, a}$ is feasible. So, $y^{o \leftarrow \emptyset, a}$ must have a strictly lower v -welfare objective than \hat{x} . Now, consider the following (in)equalities.

- $W_v(\hat{x}, \tilde{r}) > W_v(y^{o \leftarrow \emptyset, a}, \tilde{r})$ is the result derived in the paragraph above.
- $W_v(y^{o \leftarrow \emptyset, a}, \tilde{r}) = W_v(y, \tilde{r}^{o \leftarrow \emptyset, a})$, by definition of the v -welfare objective.
- $W_v(y, \tilde{r}^{o \leftarrow \emptyset, a}) \geq W_v(\hat{x}, \tilde{r}^{o \leftarrow \emptyset, a})$, by the optimality of y and the feasibility of \hat{x} .
- $W_v(\hat{x}, \tilde{r}^{o \leftarrow \emptyset, a}) = W_v(\hat{x}, \tilde{r})$, since $\hat{x}_{a\emptyset} = 0$ by assumption, and $\hat{x}_{ao} = 0$ because o is ranked below \emptyset in \tilde{r}_a .

Stringing these (in)equalities together yields $W_v(\hat{x}, \tilde{r}) > W_v(\hat{x}, \tilde{r})$, a contradiction.

To show part (b), by way of contradiction, assume that x is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $x_{a\emptyset} = 0$, but that x is not in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$. Since \hat{y} is feasible, it must be that x has a strictly lower v -welfare objective than \hat{y} . Now, consider the following (in)equalities.

- $W_v(\hat{y}, \tilde{r}^{o \leftarrow \emptyset, a}) > W_v(x, \tilde{r}^{o \leftarrow \emptyset, a})$ is the result derived in the paragraph above.
- $W_v(x, \tilde{r}^{o \leftarrow \emptyset, a}) = W_v(x, \tilde{r})$, since $x_{a\emptyset} = 0$ by assumption and $x_{ao} = 0$ since o is ranked below \emptyset in \tilde{r}_a .
- $W_v(x, \tilde{r}) \geq W_v(\hat{y}^{o \leftarrow \emptyset, a}, \tilde{r})$, by the optimality of x and the fact that $\hat{y}^{o \leftarrow \emptyset, a}$ is feasible, since $\hat{y}_{a\emptyset} = 0$.

- $W_v(\hat{y}^{o \leftrightarrow \phi, a}, \tilde{r}) = W_v(\hat{y}, \tilde{r}^{o \leftrightarrow \phi, a})$, by definition of the v -welfare objective. Stringing these (in)equalities together yields $W_v(\hat{y}, \tilde{r}^{o \leftrightarrow \phi, a}) > W_v(\hat{y}, \tilde{r}^{o \leftrightarrow \phi, a})$, a contradiction. \square

Now, I am prepared to prove [Proposition A.4](#).

Proposition A.4. Consider an agent, a , who knows the object type set is \mathcal{O} and that it contains ϕ , and fix an \mathcal{O} -conformable market realization, $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$. Let \tilde{r}_a order all object types according to the true ranking of agent a , except for ϕ .

If object type o is directly below ϕ (i.e., $\tilde{r}_{ao} = \tilde{r}_{a\phi} + 1$), then agent a weakly prefers (in the first-order stochastic sense) her expected allocation from $\tilde{r}_a^{o \leftrightarrow \phi}$ to her expected allocation from \tilde{r}_a .

Proof. To begin, I address the case where the tiebreaker gives agent a the null object type, ϕ , when she submits \tilde{r}_a . Clearly, she can't do worse to submit $\tilde{r}_a^{o \leftrightarrow \phi}$. Going forward, I will assume that when agent a submits \tilde{r}_a , the tiebreaker gives her an object type o' that is ranked higher than ϕ in \tilde{r}_a . There are two cases to consider.

There is some $y \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$ such that $y_{a\phi} > 0$. Then, by [item \(a\)](#), there must be $y' \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$ such that $y'_{a\phi} = 1$. Applying [Lemma A.4](#) with $o' = \phi$ gives $x_{a\phi} \geq 1$ for any $x \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$, since $x_{ao} = 0$ (o is unacceptable to a under \tilde{r}_a). Then, for any $x \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$, it must also be that $x_{a\phi} = 1$, which contradicts the tiebreaker giving object type o' to agent a .

For all $y \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$, it is true that $y_{a\phi} = 0$. Here, I use [Lemma A.5](#). The assignment that gives o' to agent a when she submits \tilde{r}_a plays the role of \hat{x} and any $y \in RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$ can play the role of \hat{y} . So, by part (b) of [Lemma A.5](#), \hat{x} is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$.

Now, I will show that \hat{x} is selected when the tiebreaker procedure is applied to $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$. The argument is by induction on the steps of the tiebreaker process (Equation ??) for the sets $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$. Let those steps be \mathcal{X}_t and \mathcal{Y}_t , respectively, and define condition \mathcal{C}_t to be that both $y \in \mathcal{Y}_t \Rightarrow y^{o \leftrightarrow \phi, a} \in \mathcal{X}_t$ and $\hat{x} \in \mathcal{Y}_t$ hold.

Base step (\mathcal{C}_0). Since \mathcal{Y}_0 is the set of deterministic assignments in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \phi, a}, q)$ and \mathcal{X}_0 is the set of deterministic assignments in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$, part (a) of [Lemma A.5](#) shows that for every $y \in \mathcal{Y}_0$, the assignment $y^{o \leftrightarrow \phi, a}$ is in \mathcal{X}_0 . As established above, part (b) of [Lemma A.5](#) shows that $\hat{x} \in \mathcal{Y}_0$.

Induction step ($\mathcal{C}_t \Rightarrow \mathcal{C}_{t+1}$). By way of contradiction, assume that \mathcal{C}_t holds and that \mathcal{C}_{t+1} is violated.

First, I show that \hat{x} must be in \mathcal{Y}_{t+1} . Otherwise, there must be some $z \in \mathcal{Y}_t$ that agent $\pi(t+1)$ prefers to \hat{x} . Then, by condition \mathcal{C}_t ,

it must be that $z^{o \leftrightarrow \emptyset, a}$ is in \mathcal{X}_t . But, agent $\pi(t+1)$ would also prefer $z^{o \leftrightarrow \emptyset, a}$ to \hat{x} , which means that \hat{x} is not in \mathcal{X}_t , a contradiction.

Then, the only way for \mathcal{C}_{t+1} to be violated is for there to be some $y \in \mathcal{Y}_{t+1}$ such that $y^{o \leftrightarrow \emptyset, a}$ is not in \mathcal{X}_{t+1} . But $y^{o \leftrightarrow \emptyset, a}$ gives agent $\pi(t+1)$ the same object type as y , which gives agent $\pi(t+1)$ the same object type as \hat{x} , since both y and \hat{x} are in \mathcal{Y}_{t+1} . So, for $y^{o \leftrightarrow \emptyset, a}$ to not be in \mathcal{X}_{t+1} , the assignments in \mathcal{X}_{t+1} would have to give agent $\pi(t+1)$ an object type she prefers to what she gets in \hat{x} . So, \hat{x} can't be in \mathcal{X}_{t+1} , a contradiction.

So, I have established that \hat{x} is selected when the tiebreaker procedure is applied to $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a})$. So, agent a gets o' regardless of whether she submits \tilde{r}_a or $\tilde{r}_a^{o \leftrightarrow \emptyset}$.

Hence, for any market realization, agent a receives a weakly preferred assignment when she submits $\tilde{r}_a^{o \leftrightarrow \emptyset}$ rather than \tilde{r}_a . This implies that the allocation she gets from submitting $\tilde{r}_a^{o \leftrightarrow \emptyset}$ is weakly preferred (in the first-order stochastic sense) to the allocation she gets from submitting \tilde{r}_a , as I sought to prove. \square

I can now prove [Proposition 13](#).

Proposition 13. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Further, let \mathcal{O} include the null object type, \emptyset , and let the agent have $(\mathcal{O} \setminus \{\emptyset\})$ -symmetric beliefs.*

Then, the agent weakly prefers (in the first-order stochastic sense) her expected allocation from reporting the full extension to her expected allocation from any other report.

Proof. [Proposition 11](#) shows that for any report, either truth-telling, a truncation, or an extension yields weakly greater expected utility. And [Proposition A.4](#) shows that swapping \emptyset with the object type directly below it is a weak improvement as well. The only strategy left with no improvement is full extension. \square

To show that truncations yields allocations that are stochastically preferred when the rank function is defined by [Equation 5](#), I will need an analog of [Lemma A.5](#) that works for the setup described in [Section 6.3.3](#).

Lemma A.6. *Consider an agent, a , who knows the object type set is \mathcal{O} and that it contains \emptyset , and fix an \mathcal{O} -conformable market realization, $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$. Let o be the object type directly below \emptyset in some strict ordering, $>_a$, that is, $\emptyset >_a o$ and there is no o' such that $\emptyset >_a o' >_a o$. Using the rank function ρ defined in [Equation 5](#), define $\tilde{r}_a = \rho(>_a)$, and (contrary to previous notation), also define $\tilde{r}_a^{o \leftrightarrow \emptyset} = \rho(>_a^{o \leftrightarrow \emptyset})$.*

If there is an assignment \hat{x} in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ such that $\hat{x}_{a0} = 0$, then the following statements hold.

- (a) *Any assignment in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a}, q)$ is also in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$.*

(b) Any assignment x in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ such that $x_{ao} = 0$ is also in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a}, q)$.

Proof. To prove part (a), by way of contradiction, assume some assignment y that is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a}, q)$ but not in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$. Since y is feasible regardless of the submitted rankings, it must be that \hat{x} has a strictly higher v -welfare objective. Consider the following list of (in)equalities.

- $W_v(\hat{x}, \tilde{r}) > W_v(y, \tilde{r})$ summarizes the discussion in the paragraph above.
- $W_v(y, \tilde{r}) = W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a})$, since $\tilde{r}_a^{o \leftrightarrow \emptyset}$ ranks o below \emptyset , which means $y_{ao} = 0$ (recall that in the context of [Section 6.3.3](#), \emptyset is treated the same by the rank-value mechanism, regardless of where it lies in the submitted preference).
- $W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a}) \geq W_v(\hat{x}, \tilde{r}^{o \leftrightarrow \emptyset, a})$ by the optimality of y and the feasibility of \hat{x} .
- $W_v(\hat{x}, \tilde{r}^{o \leftrightarrow \emptyset, a}) = W_v(\hat{x}, \tilde{r})$ since $\hat{x}_{ao} = 0$.

Stringing expressions together yields $W_v(\hat{x}, \tilde{r}) > W_v(\hat{x}, \tilde{r})$, a contradiction.

To prove part (b), by way of contradiction, assume x is in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $x_{ao} = 0$, but x is not in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a}, q)$. Let y be some assignment in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftrightarrow \emptyset, a}, q)$. It must be that y has a strictly larger v -welfare objective. Consider the following (in)equalities.

- $W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a}) > W_v(x, \tilde{r}^{o \leftrightarrow \emptyset, a})$ by the discussion in the preceding paragraph.
- $W_v(x, \tilde{r}^{o \leftrightarrow \emptyset, a}) = W_v(x, \tilde{r})$, since $x_{ao} = 0$, by assumption.
- $W_v(x, \tilde{r}) \geq W_v(y, \tilde{r})$ by the feasibility of y and the optimality of x .
- $W_v(y, \tilde{r}) = W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a})$ by the fact that $y_{ao} = 0$, since agent a ranks o below \emptyset in $\tilde{r}^{o \leftrightarrow \emptyset, a}$.

Stringing expressions together yields $W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a}) > W_v(y, \tilde{r}^{o \leftrightarrow \emptyset, a})$, a contradiction. \square

Before proving [Proposition 14](#), I will need one more lemma, an analog of [Proposition A.4](#) adapted to the setting of [Section 6.3.3](#).

Proposition A.5. Consider an agent, a , who knows the object type set is \mathcal{O} and that it contains \emptyset , and fix an \mathcal{O} -conformable market realization, $(b_\pi \circ RV_v, \mathcal{A}_{-a}, \tilde{r}_{-a}, q)$. Assume that when agent a submits a strict preference, the ranking submitted on her behalf will be computed using the rank function ρ , as defined in [Equation 5](#).

Let $>_a$ place all object types in the order of the true ranking of agent a , except for the null object type, \emptyset , and let o be the object type directly above \emptyset , that is, $o >_a \emptyset$ and there is no o' such that $o >_a o' >_a \emptyset$.

If object type o is truly unacceptable, then agent a weakly prefers (in the first-order stochastic sense) her expected allocation from $>_a^{o \leftrightarrow \emptyset}$ to her expected allocation from $>_a$.

Proof. Using the rank function, ρ , defined in [Equation 5](#), define $\tilde{r}_a = \rho(>_a)$, and (contrary to previous notation), also define $\tilde{r}_a^{o \leftrightarrow \emptyset} = \rho(>_a^{o \leftrightarrow \emptyset})$.

Now, let o' be the object type that the tiebreaker gives to agent a when she submits r_a . There are two cases to consider.

$o' \in \{o, \emptyset\}$. Clearly, agent a can't do worse to submit $>_a^{o \leftarrow \emptyset}$.

o' is ranked higher than o in \tilde{r}_a . Let \hat{x} denote the assignment chosen by the tiebreaker that gives agent a the object type o' . I will use [Lemma A.6](#) to show that \hat{x} is also selected by the tiebreaker procedure on $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$. The argument is by induction on the steps of the tiebreaker process (Equation ??) for the sets $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$ and $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$. Let those steps be \mathcal{X}_t and \mathcal{Y}_t , respectively, and define condition \mathcal{C}_t to be that both $y \in \mathcal{Y}_t \subseteq \mathcal{X}_t$ and $\hat{x} \in \mathcal{Y}_t$ hold.

Base step (\mathcal{C}_0). Since \mathcal{Y}_0 is the set of deterministic assignments in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$ and \mathcal{X}_0 is the set of deterministic assignments in $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}, q)$, part (a) of [Lemma A.5](#) shows that for every $\mathcal{Y}_0 \subseteq \mathcal{X}_0$. And, part (b) of [Lemma A.5](#) shows that $\hat{x} \in \mathcal{Y}_0$.

Induction step ($\mathcal{C}_t \Rightarrow \mathcal{C}_{t+1}$). By way of contradiction, assume that \mathcal{C}_t holds and that \mathcal{C}_{t+1} is violated.

First, I show that \hat{x} must be in \mathcal{Y}_{t+1} . Otherwise, there must be some $z \in \mathcal{Y}_t$ that agent $\pi(t+1)$ prefers to \hat{x} . Then, by condition \mathcal{C}_t , it must be that z is in \mathcal{X}_t as well. But, this would mean that \hat{x} is not in \mathcal{X}_t , a contradiction.

Then, the only way for \mathcal{C}_{t+1} to be violated is for there to be some $y \in \mathcal{Y}_{t+1}$ such that y is not in \mathcal{X}_{t+1} . But for y to be in \mathcal{Y}_{t+1} , it must give agent $\pi(t+1)$ the same object type as \hat{x} , since both y and \hat{x} are in \mathcal{Y}_{t+1} . So, for y to not be in \mathcal{X}_{t+1} , the assignments in \mathcal{X}_{t+1} would have to give agent $\pi(t+1)$ an object type she prefers to what she gets in \hat{x} . So, \hat{x} can't be in \mathcal{X}_{t+1} , a contradiction.

So, I have established that \hat{x} is selected when the tiebreaker procedure is applied to $RV_v(\mathcal{A}, \mathcal{O}, \tilde{r}^{o \leftarrow \emptyset, a}, q)$. So, agent a gets o' regardless of whether she submits $>_a$ or $>_a^{o \leftarrow \emptyset}$.

So, regardless of what the tiebreaker gives agent a when she submits $>_a$, she does weakly better to submit $>_a^{o \leftarrow \emptyset}$. Hence, her allocation is weakly better (in the first-order stochastic sense) when she submits $>_a^{o \leftarrow \emptyset}$. \square

Finally, I prove [Proposition 14](#).

Proposition 14. *Consider an agent who is facing a rank-value mechanism and who knows the set of object types is \mathcal{O} . Let \mathcal{O} contain \emptyset , and let the agent have $(\mathcal{O} \setminus \{\emptyset\})$ -symmetric beliefs. Further, assume when the agent submits a strict preference, the ranking submitted on her behalf will be computed using the rank function ρ , as defined in [Equation 5](#).*

Then, for any reported strict preference that isn't truncation or truth-telling, there exists another reported strict preference that is truncation or truth-telling such that the agent weakly prefers (in the first-order stochastic sense) her expected allocation from the latter report to her expected allocation from the former.

Proof. [Proposition 11](#) shows that for any report agent a might make, there is another report that is either truth-telling, a truncation, or an extension that yields an allocation that she weakly prefers (in the first-order stochastic sense). And, [Proposition A.5](#) shows that swapping \emptyset with the unacceptable object type directly above it is a weak improvement as well. The only strategies left are truth-telling and truncation. \square

B Details on the HBS FIELD Exercise

In this section, I will describe the details of the HBS FIELD match, as well as the details of the specific empirical exercise conducted in [Section 5](#) in the main text.

B.1 The HBS FIELD Match

The two distinctives that set the HBS match apart from standard random serial dictatorship are extra constraints on who can be matched together in a country and elicitation of weak preferences rather than strict (i.e., allowing agents to report indifferences).

The extra constraints were of two varieties. First, students who already have significant experience in a country are not allowed to return to that country. Second, it must be possible to split the students assigned to a country into teams of either five or six where each student is on exactly one team; each team has 0, 2, or 3 women; and no team contains more than 2 members from the same section.^{6a} An important property of these constraints is that they all depend solely on verifiable information. As such, there is no concern about truthful revelation where they are concerned. Let \mathcal{X} denote the set of assignments that meet all of these constraints, along with the more standard constraints mentioned in the main text (i.e., each student is assigned to one country and the number of students assigned to a country does not exceed its capacity). I will assume \mathcal{X} is nonempty.

A weak preference can be thought of as a strict preference over the equivalence classes induced by the indifference relation. Mathematically, define the **indifference class of object o** to be $[o] \equiv \{o' : o' \sim o\}$, and let an indifference class, $[o]$, be ranked k th if there are $k - 1$ indifference classes whose objects are strictly preferred to those in $[o]$. A **rank guarantee of k for agent a** means that agent a must be assigned a country in her k th-ranked indifference class.

The procedure used in the HBS FIELD match is formally defined by [Algorithm B.1](#). Essentially, for the k th MBA in the order, it finds the best rank guarantee it can give her while respecting the rank guarantees given to the preceding MBAs and ensuring that all other constraints are being obeyed. This mechanism is transparently strategyproof, since the preference submitted by an agent has no bearing on the opportunity set she will face when it is her turn in the order. She does best to truth-tell, since the algorithm will choose among those assignments according to whatever preference she submits.

6a. HBS students are divided into 10 sections. MBAs in the same section tend to take classes and socialize together.

INPUT: the set of valid assignments, \mathcal{X} , and a deterministic ordering of the agents, π ;

$\mathcal{X}_1 \leftarrow \mathcal{X}$;

for i from 1 to $|\mathcal{A}|$ do

$k_i \leftarrow$ the rank of the indifference class holding the most preferred allocation that agent $\pi(i)$ can get from some assignment in \mathcal{X}_i ;

$\mathcal{X}_{i+1} \leftarrow$ the assignments in \mathcal{X}_i that give agent $\pi(i)$ a rank guarantee of k_i ;

end

OUTPUT: one of the deterministic assignments in $\mathcal{X}_{|\mathcal{A}|+1}$;

Algorithm B.1: The HBS FIELD Algorithm

B.2 The Empirical Exercise

In the empirical exercise, I abstract away from the complications mentioned in the previous subsection. First, I ignore the constraints, except for those mentioned in the main text (i.e., each student is assigned to one country and the number of students assigned to a country does not exceed its capacity). As a practical matter, the omitted constraints mattered little in the actual HBS match, since there weren't many exclusions due to previous experience and given a large enough set of students, the team splits are usually attainable.

The second mode of abstraction concerns turning the submitted weak preferences into strict preferences. I accomplished this via two steps. In the first, I added all countries excluded due to previous experience to an indifference class at the bottom of the submitted preferences, so that every agent ranks all countries. In the second, I randomly broke the indifferences within each indifference class. More precisely, each indifference class in each MBA's ranking was replaced by a strict ordering of its component countries that was independently drawn from the uniform distribution over such strict orderings.

To summarize, the empirical exercise ran 10,000 iterations of the following steps: 1) form a strict preference for each agent according to the method described in the previous paragraph, 2) draw an ordering of the agents from the uniform distribution over such orderings, 3) run serial dictatorship with respect to that ordering, as well as the rank-value mechanisms described in the main text.^{7a} Finally, I average the rank distribution across all 10,000 bootstraps.

7a. The valuation vectors used in the counterfactuals are plotted in [Figure B.1](#).

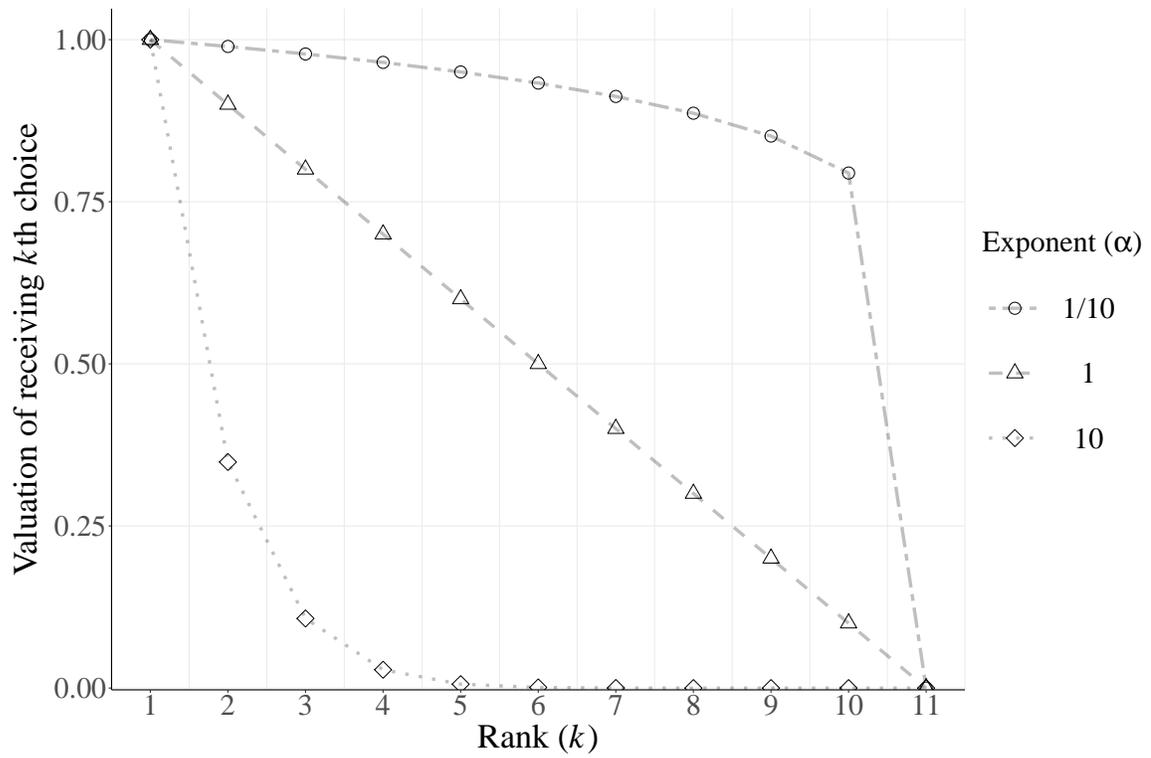


Figure B.1: Valuation Vectors Used in the Empirical Exercise

NOTES: In the legend, the numbers correspond to the exponent (α) used by the valuation formula listed in Equation 4 in the main text. All points overlap when rank is either 1 or 11, since the formula in Equation 4 in the main text is independent of the (positive) exponent at those ranks.