PICK-AN-OBJECT MECHANISMS

INÁCIO BÓ AND RUSTAMDJAN HAKIMOV

Abstract. We introduce a new family of mechanisms for one-sided matching markets, denoted pick-an-object (PAO) mechanisms. When implementing an allocation rule via PAO, agents are asked to pick an object from individualized menus. These choices may be rejected later on, and these agents are presented with new menus. When the procedure ends, agents are assigned the last object they picked. We characterize the allocation rules that can be sequentialized by PAO mechanisms, as well as the ones that can be implemented in a robust truthful equilibrium. We justify the use of PAO as opposed to direct mechanisms by showing that its equilibrium behavior is closely related to the one in obviously strategy-proof (OSP) mechanisms, but PAO-implements commonly used rules, such as Gale-Shapley DA and top trading cycles, which are not OSP-implementable. We run laboratory experiments comparing truthful behavior when using PAO, OSP, and direct mechanisms to implement different rules. These indicate that agents are more likely to behave in line with the theoretical prediction under PAO and OSP implementations than their direct counterparts.

JEL classification: C78, C73, D78, D82.
Keywords: Market Design, Matching, Sequential Mechanisms, Experiments, obvious strategy-proofness.

1. Introduction

The literature of market design, and its applications, has grown and evolved greatly over recent years. Even if we restrict our attention to the design of centralized matching markets, the instances in which theoretical and empirical contributions have influenced the way resources are allocated seem to be continuously expanding. Examples include the design of school choice procedures (Abdulkadiroğlu and Sönmez, 2003), centralized college admissions (Balinski and Sönmez, 1999), matching of resident doctors to hospitals (Roth and Peranson, 1999), organs to patients (Roth *et al.*, 2004), refugees to localities (Jones and Teytelboym, 2017), free appointments for services (Hakimov *et al.*, 2021a), and more. By carefully choosing how to determine these allocations as a function of information such

Date: August, 2021.
Bó: Southwestern University of Finance and Economics, China Center for Behavioral Economics and Finance, Chengdu, China; website: http://www.inaciobo.com; e-mail: inaciog@gmail.com.
Hakimov: Department of Economics, University of Lausanne, Switzerland and WZB Social Center Berlin, Germany. e-mail: rustamdjan.hakimov@unil.ch.
We thank Samson Alva, Lars Ehlers, Dorothea Kübler, Andrew Mackenzie, Alexander Nesterov, Marek Pycia, Madhav Raghavan, Peter Troyan, and Andrew Schotter for helpful comments. We thank Jenny Rontganger for copy editing and Nina Bonge for support in running the experiments. Hakimov acknowledges financial support from the Swiss National Science Foundation (project #100018_189152).

as participants' preferences, priority structures, and fairness concerns, these procedures can lead to allocations that satisfy certain desirable properties.

One crucial challenge faced by the designer of these mechanisms is that information needed to determine desired allocation is known by the participants but not by the designer. These are, often, their preferences over outcomes, but may also include other relevant information, such as their socioeconomic status.[1] In real-life applications, this issue is usually solved by the combination of two tools: dominant strategy implementation and the revelation principle. These tools guarantee that if the designer wants the strategic simplicity provided by dominant strategy implementation, it suffices to consider revelation mechanisms, in which participants are simply asked to report their private information, and they can safely be truthful when doing so. In fact, the vast majority of the literature focuses solely on direct revelation, and often strategy-proof, mechanisms.

However, recent experimental and empirical evidence has raised concerns about the ability of market participants to understand the incentive properties of these mechanisms. Many participants use dominated strategies, thus distorting the market allocations. This phenomenon was documented in many laboratory experiments[2] and in empirical papers.[3] These present a new challenge: is there an alternative to the implementation of allocation rules via strategy-proof direct mechanisms that would result in behavior that is more often in line with the theoretical predictions?

One recent and celebrated attempt to formalize an alternative to strategy-proofness that accounts for the extent to which participants can easily understand the incentives induced by mechanisms was the concept of obvious strategy-proofness (OSP). "A strategy is obviously dominant if, for any deviation, at any information set where both strategies first diverge, the best outcome under the deviation is no better than the worst outcome under the dominant strategy, and a mechanism is obviously strategy-proof (OSP) if it has an equilibrium in obviously dominant strategies" (Li, 2017). OSP is, therefore, a refinement of the notion of strategy-proofness, in that obvious dominance implies weak dominance. This concept could help explain why, for example, laboratory experiments indicate that agents are more likely to bid truthfully under a clock auction than under a sealed-bid second-price auction (Kagel *et al.*, 1987). While both implement the same rule in truthful dominant strategies, the former is also obviously dominant as opposed to the latter.[4]

One important shortcoming of OSP, especially for practical purposes, is that it is a very restrictive concept. Rules that are commonly considered for object allocation problems, such as top trading cycles, and stable rules, cannot be implemented via OSP mechanisms

[1]See Aygün and Bó (2021).
[2]See Hakimov and Kübler (2021) for an extensive survey of the experimental matching literature.
[3]See our related literature section.
[4]The author also presents the results of laboratory experiments comparing behavior and outcomes under strategy-proof and OSP implementations of the random serial dictatorship rule, obtaining similar results. Note, however, that Breitmoser and Schweighofer-Kodritsch (2021) raise questions as to whether the difference can be attributed to obvious strategy-proofness.

(Li, 2017; Ashlagi and Gonczarowski, 2018). This, therefore, leaves a large set of problems without this kind of behavioral guidance. Another concern is raised in a a recent paper by Pycia and Troyan (2019), who characterize all OSP mechanisms in the domain of object allocation and show that for some OSP strategies participants need to have perfect foresight to correctly predict feasible actions and outcomes, which can be demanding. To address this concern, the authors also define an even stronger concept, strong OSP, that does not require foresight from participants but leaves the designer with an essentially unique mechanism—sequential serial dictatorship (SD). In sequential SD, participants face a menu of objects in order of priority and simply pick their allocation. The strategy of choosing of the most-preferred object on the menu is *strongly obviously dominant*.

Recent laboratory experiments indicate that the dynamic implementation of the deferred-acceptance rule (DA), in which the equilibrium behavior also consists of choosing the most-preferred object from menus, leads to higher rates of truthful behavior than its standard direct revelation counterpart (Bó and Hakimov, 2020; Klijn *et al.*, 2019). The results are especially surprising given that truthful behavior is an equilibrium involving non-dominant strategies in dynamic DA, while direct DA is strategy-proof.

If one considers that the main driver behind the behavior more in line with the theory in OSP mechanisms is the fact that strategies are obviously dominant, then the forces behind the experimental results in Li (2017), Bó and Hakimov (2020), and Klijn *et al.* (2019) would be unrelated, because in the latter the equilibrium strategy is not even dominant. In this paper, we conjecture that the main driver behind the observed behavior more in line with the theoretical predictions is the simple mechanics of the equilibrium strategy, in which agents "pick" the object they would like to have from a menu, as opposed to submitting a ranking of objects representing their preferences. This would provide a unified and alternative explanation to these experimental results. Based on this, we introduce a class of sequential revelation mechanisms that implement various object allocation rules via an equilibrium behavior with closely related mechanics. We denote them *pick-an-object mechanisms* (PAO).

In a PAO mechanism, agents are asked to pick an object from individualized menus. These choices may be rejected later on, and agents are then presented with new menus containing strict subsets of the previous menus from which the previous choices have been redacted. When the procedure ends, agents are assigned the last object they picked, if any. A PAO mechanism "sequentializes" an allocation rule if it always results in the unique allocation consistent with preference profiles that could rationalize the choices made by the agents. Therefore, if agents simply choose their most-preferred object when given a menu, then the object they hold once the procedure ends is the one that the allocation rule determines given their true preferences. Notice, therefore, that truthful equilibrium behaviors in OSP and PAO mechanisms are closely related. While in the former, it can be expressed as "Wait until you can pick your best feasible object;" in the latter, it is "Pick your best feasible object and wait to see if you can keep it."

The simple mechanics involved in a PAO mechanism induces a trade-off: because information about an agent's preferences can only be obtained through choices from menus, obtaining more information about her preferences requires ruling out her last choice. This, in turn, restricts the set of allocation rules that can be "sequentialized" via PAO mechanisms.[5] We characterize such rules (Theorem 1), via a new property that we denote *monotonic discoverability*. Many familiar rules, such as the Gale-Shapley DA and top trading cycles, satisfy it (Proposition 1). We characterize the rules that can be implemented in truthful strategies via a robust equilibrium (robust ordinal perfect Bayesian equilibrium) as being those that are strategy-proof and satisfy monotonic discoverability (Theorem 2). Finally, we show that every non-bossy OSP-implementable rule is implementable in weakly dominant strategies via PAO mechanisms (Theorem 3).

Our justification for the use of PAO mechanisms is non-standard, in the sense that the game-theoretical incentive properties of the PAO mechanisms are weaker than the alternatives, that is, strategy-proof direct mechanisms and OSP mechanisms, but we test our conjecture via laboratory experiments.

We test two allocation rules, the top trading cycles (TTC) and serial dictatorship (SD), and we construct three treatments for each: direct revelation implementation, PAO implementation, and OSP implementation. That is, for each one of these rules, we ran each of the three different mechanisms for implementing them. We changed the mechanisms implementing the rules between-subjects and the main goal of the experiment is to compare performance of PAO mechanisms relative to the direct and OSP (when possible) counterparts in various environments (in our case rules). We changed the allocation rules within-subjects out of practical considerations and because we are interested in comparative statics between the types of mechanisms, but not the rules. Because TTC is only OSP-implementable for certain "acyclic" priority structures (Troyan, 2019), we split the TTC environments into cyclic and acyclic priority structures, with the former having only direct and PAO implementations.

We find that, in fact, OSP implementations lead to higher truth-telling rates across the board. When the OSP implementation exists (i.e., except for TTC with cyclic priority structures), OSP outperforms both PAO and direct implementations in terms of truthful preference revelation. When comparing the direct implementation of TTC vs. the PAO mechanism, the experiments show that the PAO implementation leads to a higher proportion of subjects following truthful equilibrium strategies for the TTC rule and no difference for the SD rule.[6] As for efficiency, PAO mechanisms lead to significantly higher efficiency than direct mechanisms for both TTC and SD. OSP mechanisms improve efficiency over the direct ones, but there is no significant difference to PAO mechanisms. Thus, despite a higher proportion of truthful strategies, OSP mechanisms do not improve

[5]We provide a simple motivating example in section 2.1.

[6]The experiments reported in Bó and Hakimov (2020) complement these with a comparison between the direct revelation Gale-Shapley DA with the iterative deferred acceptance mechanism, which is its PAO implementation. The results are in line with the ones that we present here: the PAO implementation of DA results in a higher proportion of truth-telling than its direct counterpart.

efficiency relative to PAO ones. This is because the deviations from the truthful strategy are more likely to be payoff irrelevant in OSP than in PAO mechanisms.

When we look more closely at the results for the OSP implementation of TTC, however, we see a big difference in the rates of truthful behavior, depending on the nature of the obviously dominant strategy for a subject. We employ the characterization of Pycia and Troyan (2019), who show that every OSP-implementable mechanism is equivalent to a "millipede game." In millipede games, at every decision node each player has a choice between leaving with an object among those in a given menu ("clinching action"), and at most one "passing action." Unlike the clinching actions—in which agents simply choose their allocation—the passing actions require foresight to correctly predict feasible options, which can be demanding. In our experiment, in the OSP implementation of the acyclic TTC rule, when the obviously dominant strategy consists of simply picking the most-preferred object—the "clinching action"—the rate of truthful behavior is 93%. But when it requires some degree of "foresight," in that it involves the "passing action," the rate of truthful behavior is only 56%. This result strongly supports the strong OSP concept of Pycia and Troyan (2019), and its myopic "picking" equilibrium as being a better predictor of behavior than OSP in general.

Thus far, we motivated the PAO mechanisms only by the attractiveness of straightforward strategies empirically, which was confirmed by our experiments. However, PAO mechanisms have other attractive features, which, although not directly evaluated in this paper, are important in practical applications. First, they improve the acquisition of information about the options available during the execution of the procedure, by requiring coarser information about preferences and limiting the number of options available between steps (Grenet *et al.*, 2019; Hakimov *et al.*, 2021b). Second, by allowing the participants to *experience* the steps involved in the production of the final allocation, they can be perceived as more transparent (Hakimov and Raghavan, 2020a). Finally, they can make the equilibrium strategy more feasible in markets with a very large number of options, such as nationwide college admissions,[7] when compared to direct mechanisms. This is because in most practical cases, designers constrain the length of a rank-order list, as ranking even 100 options seems to be a very hard task, while the number of choices that students have to make in a PAO implementation of DA, for example, is typically much smaller than the number of alternatives that should be ranked (Bó and Hakimov, 2016). This, of course, comes at the cost of a longer time for the mechanism to run, which is an important practical consideration.[8] Finally, the use of dynamic college admission mechanisms seem to be trending in recent reforms of college admissions, like in France, Inner-Mongolia, Germany, and Tunisia (Bó and Hakimov, 2016; Gong and Liang, 2016; Luflade, 2018).

[7]During university admissions in China and Brazil, for example, students face thousands of programs and universities (Gong and Liang, 2016; Bó and Hakimov, 2016).

[8]In college admissions in France, which runs a mechanism where students dynamically receive offers from colleges, the deadline for decision ranges from 5 days at the start of the procedure to 1 day towards the end. The system has been in place since 2018.

To sum, we show that the PAO environment has significant benefits over its direct counterpart in the general decision environment. However, when there is the option of using OSP mechanism, our experiments suggest using them. Because many allocation rules used in real-life allocation problems are not OSP-implementable, but are PAO-implementable, we interpret our experimental results as support for the choice of PAO mechanisms over their direct revelation counterparts in these cases.

Related Literature. In addition to the studies mentioned in the introduction, this paper is mainly related to two literature strands: the design of sequential allocation mechanisms and their theoretical properties and the behavioral and experimental aspects of market design.

From the theoretical perspective, perhaps the closest paper to ours is Mackenzie and Zhou (2020). They consider the family of *menu mechanisms*, which are also sequential revelation mechanisms in which participants are asked to choose from menus of their possible assignments. As in the case of PAO mechanisms, they focus on those in which an agent can never select an assignment twice. Unlike PAO mechanisms, however, the definition of menu mechanisms does not imply a restriction on the set of allocation rules that can be sequentialized because an agent's assignment does not have to be the last choice of an agent and may, in fact, be an object that was not chosen at all. Despite considering this more general setup, they show that strategy-proof rules can be implemented in a truthful ex-post perfect equilibrium, a result similar to our Theorem 2.

Another closely related paper is Börgers and Li (2019). As in our case, they are concerned about a notion of simplicity in mechanisms. They define the class of "strategically simple" mechanisms, which are those in which an agent's optimal strategy depends only on first-order beliefs about preferences and rationality. Like the rules implemented in truthful equilibria in PAO mechanisms, these include all dominant strategy mechanisms and extend to others. It is worth noting, however, that PAO mechanisms are not necessarily strategically simple.

Other papers have also considered sequential versions of allocation rules, such as multi-unit auctions (Ausubel, 2004, 2006), stable matchings (Bó and Hakimov, 2016; Kawase and Bando, 2021; Haeringer and Iehle, 2019) and more general allocations (Schummer and Velez, 2021). Moreover, there is a growing literature evaluating sequential mechanisms used in the field (Gong and Liang, 2016; Grenet *et al.*, 2019; Veski *et al.*, 2017; Dur *et al.*, 2018). Other recent papers, such as Akbarpour and Li (2020) and Hakimov and Raghavan (2020b) show that the use of sequential mechanisms can also be explained by their transparency and credibility characteristics: the experience that participants have when interacting with these mechanisms can convey information that helps them to be sure that the allocation is produced by following the rules.

We also provide contributions to the literature documenting dominated behavior in matching mechanisms. Shorrer and Sóvágó (2018); Rees-Jones (2018); Hassidim *et al.* (2016) and Artemov *et al.* (2017) document dominated strategies being played in real-life centralized allocation processes. These are in line with laboratory experiments that also evaluate truthful behavior in strategy-proof mechanisms relative to dynamic mechanisms (Echenique *et al.*, 2016; Bó and Hakimov, 2020; Breitmoser and Schweighofer-Kodritsch, 2021; Klijn *et al.*, 2019), and truthful behavior in TTC and SD (Chen and Sönmez, 2002, 2006; Pais and Pintér, 2008; Guillen and Hakimov, 2017, 2018; Hakimov and Kesten, 2018).

2. Model

Let $A = \{a_1, a_2, \ldots, a_n\}$ be a finite set of **agents**[9] and $O = \{o_1, o_2, \ldots, o_m\} \cup \{\emptyset\}$ be a set of **object types**, where $\emptyset$ is the **null object**. Each agent a has strict **preferences** P_a over the set O. Given P_a we express the induced weak preference by R_a. That is, oR_ao' if oP_ao' or $o = o'$. We abuse notation and P_a may represent its binary relation (ex: oP_ao') or a tuple of elements of O, for example, $P_a = (o, o', \emptyset, o'')$, which implies $oP_ao'P_a\emptyset P_ao''$. We also often treat tuples of distinct elements of O as sets, if that does not create any ambiguity. For example, we may say that $\gamma = (o_1, o_2, o_3)$ and $\{o_2\} \subset \gamma$. Denote by $\mathbb{P}$ the set of all strict preferences over O. An object $o \in O$ is **acceptable** to agent $a \in A$ if $oP_a\emptyset$. A **preference profile** is a list $P = (P_{a_1}, P_{a_2}, \ldots, P_{a_n})$. We denote by P^{-a} the set of all preferences in P except for P_a. A **problem** is a triple $\langle A, P, O\rangle$. Let $\mathcal{P} = \mathbb{P}^n$ be the set of all preference profiles. An **allocation** is a function $\mu : A \to O$.[10] For a given allocation μ, we say that **agent a's assignment** under μ is $\mu(a)$. Let $\mathcal{M}$ be the set of all allocations. A random allocation is a probability distribution over $\mathcal{M}$. A **rule** is a function $\varphi : \mathcal{P} \to \mathcal{M}$. Denote by $\varphi_a(P) = \varphi(P)(a)$. A rule φ is **individually rational** if, for any $P \in \mathcal{P}$ and $a \in A$, $\varphi_a(P)\, R_a \emptyset$.

A rule is **strategy-proof** if for every agent $a \in A$, $P \in \mathcal{P}$, and $P' \in \mathbb{P}$, $\varphi_a(P_a, P^{-a})\, R_a \varphi_a(P', P^{-a})$.

Define a **choice history** h as a sequence of tuples $((\Omega_1, \omega_1), (\Omega_2, \omega_2), \ldots)$, where for every i, $\Omega_i \subseteq O$ and $\omega_i \in \Omega_i$. That is, a choice history is a sequence of sets of object types and elements of those sets. For example:

$$((\{o_1, o_2, o_3, o_4, o_5\}, o_2), (\{o_3, o_4, o_5\}, o_5))$$

Because this will be used often in what follows, we denote by $\overrightarrow{h}$ the last choice in h. That is, if $h = ((\Omega_1, \omega_1), \ldots, (\Omega_k, \omega_k))$, $\overrightarrow{h} = \omega_k$. We say that h is a **continuation history** of h' if all the tuples in h' are also in h. We denote by H the set of all choice histories, including the empty choice history, represented by $h^{\emptyset}$. We say that a preference P_a is **consistent with the choice history** h if for every $(\Omega_i, \omega_i) \in h$ and $o \in \Omega_i$, $\omega_i R_a o$. We denote by $P(h)$ the set of all preferences that are consistent with h. We can also say that a choice history is consistent with a preference using inverse reasoning, and denote by $h(P_i)$ the set of all choice histories that are consistent with P_i.

[9]To simplify notation, an agent a_i might sometimes be denoted by her index i.

[10]Notice that while this is a model of discrete object allocation, there is no explicit notion of feasibility considering capacities. Feasibility is "encoded" in the allocation rules themselves: if an allocation is in the image of the rule, then it is feasible.

Denote by **collective history** h^A a list of n choice histories: $(h_1, h_2, \ldots, h_n)$. We denote by H^A the set of all collective histories, by $h^{A-\emptyset}$ the collective history consisting of n empty choice histories, and by h_i^A the i-th element in h^A. We also say that h^A is a **continuation collective history** of $h^{A'}$ if each choice history in h^A is a continuation history of its related history in $h^{A'}$. A preference profile $P = (P_{a_1}, P_{a_2}, \ldots, P_{a_n})$ is **consistent with the collective history** h^A if for every i, P_{a_i} is consistent with h_i^A.[11] We denote by $P\left(h^A\right)$ the set of preference profiles that are consistent with h^A. We can also say that a collective history is consistent with a preference profile using inverse reasoning, and denote by $h^A(P)$ the set of all collective histories that are consistent with the preference profile P.

Next, we abuse notation and define $\varphi\left(h^A\right)$ to be the allocations that are consistent with the application of the rule φ, given the information about the preference profile that can be deduced from taking the revealed preference approach to the the collective history h^A:

$$\varphi\left(h^A\right) = \bigcup_{P \in P(h^A)} \varphi(P)$$

For a given collective history h^A, therefore, $\varphi\left(h^A\right)$ is a (possibly empty) subset of $\mathcal{M}$. We also define the set of **feasible assignments after** h^A **for** a_i, or $\mu_i^{\varphi}\left(h^A\right)$ as:

$$\mu_i^{\varphi}\left(h^A\right) = \bigcup_{\mu \in \varphi(h^A)} \mu(a_i)$$

Let $\Phi = \left(2^O\right)^n$. That is, Φ is the set of n-tuples with subsets of O.

A **menu function** $\mathbb{S} : H^A \to \Phi$ specifies, for each collective history, a list of menus to be given to the agents.

- $\mathbb{S}\left(h^{A-\emptyset}\right) = (\phi_1^0, \phi_2^0, \ldots, \phi_n^0)$, where for every i, ϕ_i^0 is a non-empty subset of O. These are the initial menus.
- For any collective history h^A, where $h_i^A = \left((\Omega_1^i, \omega_1^i), (\Omega_2^i, \omega_2^i), \ldots, \left(\Omega_{k_i}^i, \omega_{k_i}^i\right)\right)$,[12] $\mathbb{S}\left(h^A\right) = (\phi_1, \phi_2, \ldots, \phi_n)$, where $\phi_i \subseteq \Omega_{k_i}^i \backslash \{\omega_{k_i}^i\}$.

Denote by $\mathbb{S}^i(\cdot)$ the i-th element in $\mathbb{S}(\cdot)$ (or, when convenient, $\mathbb{S}^a(\cdot)$ to be the element in $\mathbb{S}(\cdot)$ associated with agent a). For a given menu function $\mathbb{S}$, we define the **pick-an-object mechanism** $\mathbb{S}$ as follows:

- Period $t = 1$: For every agent a_i, ask her to choose one item in $\mathbb{S}^i\left(h^{A-\emptyset}\right)$. Let agent a_i's choice be σ_i^t. Define h^{A-1} to be the collective history such that for every agent a_i, $h_i^{A-1} = ((\phi_1^0, \sigma_i^1))$.
- Period $t > 1$: Let $(\phi_1^t, \phi_2^t, \ldots, \phi_n^t) = \mathbb{S}\left(h^{A-t}\right)$.

[11] Note that this implies that a preference profile is consistent with a collective history if **all** histories are consistent with some preference. In other words, it is necessary that all histories are rationalizable.

[12] For each agent a_i, k_i represents the number of menus that she was given in h^A.

– If for all i, $\phi_i^t = \emptyset$, then the procedure stops, and outputs the allocation μ, where for each i, $\mu(a_i) = \overrightarrow{h_i^{A-t}}$.
– Otherwise, for every agent a_i, ask her to choose one item in ϕ_i^t, if the menu is non-empty.[13] Let agent a_i's choice be σ_i^t. Define $h^{A-(t+1)}$ as the collective history such that for every agent a_i who received a non-empty menu, $h_i^{A-(t+1)} = h_i^{A-t} \oplus (\phi_i^t, \sigma_i^t)$,[14] and for those with an empty menu, $h_i^{A-(t+1)} = h_i^{A-t}$.

Notice that because the menus given to each agent do not include her previous choices and never include objects that were not present in previous menus, every collective history that results from any number of periods of a PAO mechanism is consistent with a non-empty set of preference profiles. Moreover, since the agents' allocation must be their last choice and menus are subsets of previous ones, every feasible allocation for an agent must be in the first menu given at $t = 1$. When facing a PAO mechanism, one simple behavior that an agent may follow is what we call a **straightforward strategy**, which we define below.

Definition 1. An agent follows a **straightforward strategy with respect to** P if whenever presented with a menu $I \subseteq O$, she chooses the most-preferred element of I according to P.

Definition 2. A pick-an-object mechanism $\mathbb{S}$ **sequentializes** the rule φ if, for any preference profile P, the pick-an-object mechanism $\mathbb{S}$ provides menus such that when each agent a_i follows the straightforward strategy with respect to P_{a_i}, the outcome $\varphi(P)$ is produced. We say that there exists a pick-an-object mechanism that sequentializes some rule φ if there exists a menu function $\mathbb{S}$ such that a pick-an-object mechanism $\mathbb{S}$ sequentializes φ.

2.1. **Monotonic discoverability.** At first glance, it might seem like every rule can be sequentialized by some pick-an-object mechanism. After all, by asking agents to choose from menus, one can always recover as much information about their preferences as necessary to pinpoint an allocation for a given rule. The definition of PAO mechanisms, however, imposes some restrictions on the menus that can be presented to an agent, and how that relates to her assignment. In particular, objects previously chosen cannot be in future menus, and the agent is assigned to the last object she chose. These conditions result in mechanisms that have a simple and intuitive operation, but also induce a trade-off between obtaining more information about her preference and the assignment that the rule determines for her.

To see this, consider a problem where $O = \{o_1, o_2, o_3, \emptyset\}$, $A = \{a\}$, and the rule φ^* as follows:

[13] An agent receiving an empty menu represents a situation in which she is not called to make a choice. One could alternatively interpret agents receiving empty menus as "inactive" agents in this period.
[14] We use "$\oplus$" to denote concatenation. That is, for example, $((\phi,\sigma),(\phi',\sigma')) \oplus (\phi'',\sigma'') = ((\phi,\sigma),(\phi',\sigma'),(\phi'',\sigma''))$.

$$o_1 P_a o_2 P_a o_3 \; : \; \mu(a) = o_1$$
$$o_1 P'_a o_3 P'_a o_2 \; : \; \mu(a) = o_3$$
$$\text{Otherwise} \; : \; \mu(a) = \emptyset$$

Suppose that we are looking for a PAO mechanism that sequentializes φ^*, and consider first which object types should be in menu given to the agent in period $t = 1$. Clearly, object types o_1 and o_3, as well as the null object $\emptyset$ must be in the menu, since the agent's potential assignments must be in it. So there are two possibilities for the first menu: $\phi = \{o_1, o_2, o_3, \emptyset\}$ and $\phi' = \{o_1, o_3, \emptyset\}$.

What happens if the agent chooses o_1 from the first menu? If the menu was ϕ, this tells us that o_1 is the most preferred object type. This, however, does not provide enough information to narrow down to a single allocation, since this choice is consistent with the preferences P_a and P'_a above, but also with, for example, $o_1 P''_a o_3 P''_a \emptyset P''_a o_2$. More information about the agent's preferences is necessary to know which allocation to produce.

More information about the agent's preferences can only be obtained by observing her choice from another menu. As per the definition of PAO mechanisms, that new menu and all the following ones cannot have o_1 as one of the options. We reach, therefore, an **informational gridlock**: in order to know which assignment to produce, we need to know, at the very least, whether o_2 is preferred to o_3 or not. Obtaining more information "costs" eliminating o_1 as a possible assignment for a. But if her preference is P_a, the rule φ^* indicates that that should be her assignment. The same problem is present if the first menu was ϕ': if she chooses o_1 or o_3, we reach the same gridlock.[15]

We conclude, therefore, that φ^* cannot be sequentialized by a PAO mechanism. As we will show next this is because φ^* does not satisfy *monotonic discoverability*, a property that we now define.

Let μ be an allocation. Let $P = (P_{a_1}, P_{a_2}, \ldots, P_{a_n})$ be any element of $\mathcal{P}$. We define the function $\mathcal{L}(P, \mu) = \{P' \in \mathcal{P} : \forall a \in A \text{ and } o, o' \in O : o P_a o' P_a \mu(a) \iff o P'_a o' P'_a \mu(a)\}$. That is, $\mathcal{L}(P, \mu)$ contains all preference profiles that, for each agent a, agree with P_a with respect to $\mu(a)$ and all objects preferred by a to $\mu(a)$, but may differ with respect to objects that $\mu(a)$ are preferred to, with respect to P_a. We will say that $\mathcal{L}(P, \mu)$ is therefore the *lower contour set* of P at μ. We will denote each element of $\mathcal{L}(P, \mu)$ a **continuation profile of P at μ.**

Definition 3. A rule φ satisfies **monotonic discoverability** if, for any allocation μ and preference profile P, either $\varphi(P) = \mu$ or there is an agent $a^* \in A$ such that $P' \in \mathcal{L}(P, \mu) \implies \mu(a^*) \neq \varphi_{a^*}(P')$.

[15]Note that this observation does not rely on the fact that the definition of PAO mechanisms imply that the first menu must contain all the agent's feasible assignments. The type of "informational gridlock" is obtained if we simply require the agent to be assigned her last choice and menus to be such that choices are always rationalizable.

Step 1

P_{a_1}	o_1	o_2	o_3	o_4	o_5	$\emptyset$
P_{a_2}	o_3	o_1	o_4	o_3	o_5	$\emptyset$
P_{a_3}	o_2	o_5	o_4	o_1	o_3	$\emptyset$
P_{a_4}	o_5	o_4	o_2	$\emptyset$	o_1	o_3

Step 2

P_{a_1}	o_1	o_2	o_3	o_4	o_5	$\emptyset$
P_{a_2}	~~o_3~~	o_1	o_4	o_3	o_5	$\emptyset$
P_{a_3}	o_2	o_5	o_4	o_1	o_3	$\emptyset$
P_{a_4}	o_5	o_4	o_2	$\emptyset$	o_1	o_3

Step 3

P_{a_1}	~~o_1~~	o_2	o_3	o_4	o_5	$\emptyset$
P_{a_2}	~~o_3~~	o_1	o_4	o_3	o_5	$\emptyset$
P_{a_3}	o_2	o_5	o_4	o_1	o_3	$\emptyset$
P_{a_4}	~~o_5~~	o_4	o_2	$\emptyset$	o_1	o_3

Step 4

P_{a_1}	~~o_1~~	o_2	o_3	o_4	o_5	$\emptyset$
P_{a_2}	~~o_3~~	o_1	o_4	o_3	o_5	$\emptyset$
P_{a_3}	o_2	o_5	o_4	o_1	o_3	$\emptyset$
P_{a_4}	~~o_5~~	~~o_4~~	o_2	$\emptyset$	o_1	o_3

FIGURE 1. Monotonic discoverability and sequentialization

To understand the critical role that monotonic discoverability has when sequentializing rules, consider Figure 1. Suppose that we have a rule φ that satisfies monotonic discoverability, and start with the allocation μ^0 that matches each agent with her most-preferred object type, highlighted in blue in Figure 1, step 1. We can consider two cases. In the first, φ is such that for every preference profile in which agent a_i's top option is as in P, these agents are matched to their top choices. In that case, knowing those top choices already gives us enough information to determine this to be the outcome that φ maps for any continuation profile.[16] The second case is when this is not true. That is, there are continuation profiles of P at μ^0 for which φ does not map those profiles to μ^0. Let P^0 be one such profile. By monotonic discoverability, there is at least one agent who, for every continuation profile of P^0 at μ^0 (and therefore also of P at μ^0) will not be matched to her match at μ^0. Without loss of generality, let a_2 be such an agent.

Consider next the allocation μ^1, which is the same as μ^0 except that the object mapped to a_2 is the second one in her preference. That is shown in Figure 1, step 2. Here, once again, we have two cases. In the first, φ is such that for every continuation preference profile of P at μ^1,all agents are matched to their outcomes under μ^1. The second case is where there is at least one, but potentially many, such profiles in which the outcomes are different than μ^1. There is one thing we can say, however. Because the continuation profiles of P at μ^1 are also continuation profiles of P at μ^0, for none of these cases is agent a_2 matched to o_3. Suppose here, without loss of generality, that the second case is true, agents a_1 and a_4 are not matched to their outcomes under μ^1 for any continuation profile of P at μ^1. Steps 3 and 4 represent a continuation of this argument, until at step 4, the allocation μ^4 being evaluated is in fact the one that the rule φ maps to all continuation

[16]One simple example of this situation is when φ is a simple serial dictatorship, and no two agents have the same object as their most-preferred one, as in P.

profiles of P at μ^4. Notice that this process always converges to the allocation mapped by φ to the profile P, because we only move down the preference ordering of an agent when we have enough information to determine that φ rules out previously considered matches to her.

Following this process shows us two consequences of a rule satisfying monotonic discoverability. The first is that while following this monotonic process of evaluating allocations that are at each step weakly worse (from the perspective of the underlying preference profile), we end up at the allocation that φ maps to P. This fact is fundamental for a PAO mechanism to be able to match agents to their last choices without repetition.

The second is that monotonic discoverability solves the informational gridlock that is induced by a PAO mechanism; that is,we must permanently reject the last choice agent made to obtain more information about her preferences. Monotonic discoverability guarantees that, as long as we follow the monotonic process over allocations that we just described, whenever we do not have enough information to single out an allocation (that is, there are still continuation profiles for which φ produces different allocations), there is at least one agent whose last choice can be rejected and then presented with a new menu, allowing us to obtain more information about preferences, narrowing down the set of preference profiles consistent with these choices. In fact, this relationship between monotonic discoverability and PAO mechanisms is as strong as possible, as shown in Theorem 1 below.

Theorem 1. *There exists a pick-an-object mechanism that* ***sequentializes*** *an individually rational rule φ if and only if φ satisfies monotonic discoverability.*

Theorem 1 implies, therefore, that the structure of PAO mechanisms limits the rules that can be sequentialized using them. If we instead decide to obtain information about preferences via choices from menus, we would have no such restriction: one could, for example, make each agent choose one object at a time starting from the entire set O. This elicits the entire preference profile, which could then be used to determine the allocation. The participant's experience in this type of mechanism would be fundamentally different, though: the connection between their choices and assignments would be unclear, in that they might end up matched to an object chosen early in the process, before many other choices. In PAO mechanisms, there is a much more explicit connection between choices and assignments: an agent can pick what she wants and will be able to keep it unless the information provided by the other participants, through their choices, determines that she cannot keep it, in which case she is offered the chance of choosing again from a new menu with feasible options.[17]

2.2. Generalized Deferred Acceptance Procedures.

Many mechanisms used in matching are defined by algorithms that produce outcomes, as opposed to axioms or objective

[17]The specific implementation of the PAO mechanism could even allow for the market designer to credibly and truthfully explain the reason why the agent cannot keep the object, using the information that she obtained that resulted in that rejection.

functions. Two classic examples are the Gale-Shapley deferred acceptance mechanism (DA) (Gale and Shapley, 1962) and Gale's top trading cycles (TTC) (Shapley and Scarf, 1974). Their definitions describe step-by-step procedures that use agents' preference rankings (and often other information, such as priority orderings) and result in an allocation.

Generalized DA procedures is a general description of those which produce tentative allocations at each step, following the agent's preferences, until no agent has her choice rejected. They generalize DA in the sense that they determine whether the tentative allocation of an agent to an object type becomes a rejection based on the *entire* tentative allocation and set of proposals, as opposed to the tentative allocation and proposals *to the object type in question*. If we use the college admissions analogy, in a generalized DA procedure, whether a student's proposal is accepted or not may depend not only on the college she applied to (and its tentatively matched students) but also on the entire tentative assignments of students to colleges, and contemporaneous applications.

A generalized DA procedure is, therefore, defined by an *update function* $\Psi : \mathcal{M} \times \mathcal{M} \to \mathcal{M}$. In it, $\Psi(\mu, \mu')$ informs, for a "tentative" assignment μ, what will be the new tentative assignment when the new proposals are the ones represented by the assignment μ'. This has the restriction that, if $\Psi(\mu, \mu') = \mu''$, for every $a \in A$ it must be that $\mu''(a) \in \{\emptyset, \mu(a), \mu'(a)\}$. That is, the new tentative allocation must be one that for each agent, either leaves him with the same tentative allocation, unmatched, or tentatively assigned to the object type she has just proposed. Given an update function Ψ, the procedure can be described by the following algorithm:

- **Step** 1: Let μ^0 be an assignment in which, for all $a \in A$, $\mu^0(a) = \emptyset$, and μ^* be an assignment consisting of all agents matched to their top choice in P. Let, moreover, $\mu^1 = \Psi(\mu^0, \mu^*)$. We say that every agent $a \in A$ for which $\mu^*(a) \neq \emptyset$ and $\mu^1(a) = \emptyset$ was *rejected by* $\mu^*(a)$.
- **Step** $t > 1$: Construct the assignment μ^*, in which every agent a who was rejected at step $t-1$ is matched to her most-preferred object, with respect to P_a, from which she was not previously rejected. Moreover, let all other agents be matched to $\emptyset$ in μ^*. Let, moreover, $\mu^t = \Psi(\mu^{t-1}, \mu^*)$. If for every $a \in A$ it is the case that $\mu^t(a) \in \{\mu^{t-1}(a), \mu^*(a)\}$, stop the procedure and determine the assignment to be μ^t. Otherwise, go to step $t+1$.

Rules that are described by generalized DA procedures satisfy monotonic discoverability, as shown below.

Proposition 1. *If φ is described by a generalized DA procedure, then φ satisfies monotonic discoverability.*

Proposition 1 implies that many mechanisms that are commonly used and referenced in the matching literature satisfy monotonic discoverability and can therefore be sequentialized by a PAO mechanism. One can easily see that these include, for example, DA itself, TTC, and the Boston mechanism.

2.3. **A canonical Pick-an-object mechanism.** Up to this point, we described how PAO mechanisms operate, and which rules can be sequentialized with PAO mechanisms. The next natural question is: given some rule φ that satisfies monotonic discoverability, how can we design the contents of the menus given to the agents in each step, in a way that sequentializes φ?

In principle, there might be multiple PAO mechanisms that sequentialize a given rule, with variations over which agents are asked to choose from menus and the contents of those menus. We can, however, construct a **canonical PAO mechanism** for that rule. Let φ be a rule that satisfies monotonic discoverability, and define the PAO function $\mathbb{S}$, such that for every i and $h^A \in H^A$:

$$\mathbb{S}^i\left(h^A\right) = \begin{cases} \emptyset & \text{if } \left|\varphi\left(h^A\right)\right| = 1 \text{ or } \overrightarrow{h_i^A} \in \mu_i^\varphi\left(h^A\right) \\ \mu_i^\varphi\left(h^A\right) & \text{otherwise} \end{cases}$$

The proof of Theorem 1 involves showing that the PAO mechanism $\mathbb{S}$ above sequentializes the rule φ. The resulting PAO mechanism can also be easily explained:

- Step 1: Every agent i is given a menu containing $\mu_i^\varphi\left(h^\emptyset\right)$, that is, all the object types that are allotted to i in some allocation produced by the rule φ.[18]
- Step $t > 1$: Let h^A be the collective history representing the menus and choices made by the agents up to period t. There are two cases.
  - $\left|\varphi\left(h^A\right)\right| = 1$, that is, every preference profile consistent with the collective history up to step t is mapped to the same assignment by φ. In that case, by monotonic discoverability, the last object chosen by the agents is exactly that assignment, and therefore the procedure ends and agents leave with the last object they picked.
  - Otherwise, monotonic discoverability implies that the set $A^* \equiv \left\{i \in A : \overrightarrow{h_i^A} \notin \mu_i^\varphi\left(h^A\right)\right\}$ is non-empty. This set contains each agent for which, for any preference profile consistent with the collective history h^A, the assignment given by φ is different from her last chosen object type. Each one of these agents in A^* are given a menu containing the set of object types that are still feasible for them, under φ, for the preference profiles consistent with the collective history h^A.

The canonical PAO mechanism gives, therefore, for a given rule φ a simple recipe for how the menus should be constructed. For some rules the resulting mechanism is very intuitive.[19] For some other, such as the simple serial dictatorship (SD), it might be less so, and alternatively formulated PAO implementations are "simpler" to understand (see section 4).

Another fact worth noting is that when a rule is individually rational, every menu given by the canonical PAO mechanism contains the element $\emptyset$, and whenever an agent chooses

[18] In other words, the set of feasible assignments for i under φ. In a college admissions environment, for example, that would be every college that deems i acceptable.

[19] The Iterative Deferred Acceptance Mechanism (Bó and Hakimov, 2016), for example, is the canonical PAO mechanism for the DA rule.

it, her assignment is determined to be the null object—that is, she is left without any object.

Monotonic discoverability therefore gives a full characterization of the rules for which it is possible to use a PAO mechanism to "sequentialize" the process of obtaining information about the participants' preferences, and Proposition 1 gives us a family of mechanisms that satisfy that condition. This does not, however, guarantee that it will be in the participants' own interest to truthfully reveal their preferences. When that is not the case, then the outcomes produced by these PAO mechanisms may differ substantially, with respect to the agents' true preferences, with that determined by the rule being used. In other words, we also need to consider the implementation problem when using PAO mechanisms.

3. Implementation in Pick-an-Object Mechanisms

In this section we will consider the extensive-form game that is induced by PAO mechanisms, and conditions on the rules that guarantee that it is in the participants' own interest to reveal their true preferences using straightforward strategies. The equilibrium concept that we will use is that of ordinal perfect Bayesian equilibrium (OPBE), introduced in Bó and Hakimov (2016). Loosely speaking, a strategy profile is an OPBE if at every information set, following the equilibrium strategy first-order stochastically dominates any deviating strategy. It is, therefore, an ordinal version of a perfect Bayesian Nash equilibrium that will be formalized later in this section.

Straightforward strategies being an OPBE implies that choosing from menus according to the agent's true preferences first-order stochastically dominates any other strategy *regardless of past actions by the player being considered*. That is, it implies that providing truthful information (in the form of choices based on true preferences) is the best thing to do even when the information that was provided earlier was not truthful.[20]

Next, we formalize the game description, the participants' information and beliefs. To do that, we will expand the definition of a PAO mechanism to also include the information that it provides to the participants. We do this by defining, for a given mechanism, the information structure that is associated with it.

Let $\mathbb{S}$ be a menu function. We denote by $H^A_{\mathbb{S}}$ the set of all collective histories that can result from the PAO mechanism $\mathbb{S}$. That is, $H^A_{\mathbb{S}}$ contains each collective history that would result from each possible combinations of choices from all agents from the menus that are offered when using the PAO mechanism $\mathbb{S}$. Next, for each agent $a \in A$, let $\mathcal{I}^a$ be **agent a's information structure**. $\mathcal{I}^a = \{\mathcal{I}^a_1, \mathcal{I}^a_2, \dots\}$ is a partition of the set $H^A_{\mathbb{S}}$, such that for every pair of collective histories $h^A, h^{A\prime} \in H^A_{\mathbb{S}}$:

(1) If $h^A = \left(h^A_{a_1}, \dots, h^A_a, \dots, h^A_{a_n}\right)$ and $h^{A\prime} = \left(h^{A\,\prime}_{a_1}, \dots, h^{A\prime}_a, \dots, h^{A\,\prime}_{a_n}\right)$ and $h^A_a \neq h^{A\prime}_a$, then h^A and $h^{A\prime}$ must be in different elements of $\mathcal{I}^a$.

[20] This seems, at first sight, stronger than strategy-proofness, in which all that is required is that reporting the true type, and just that, is always a best-response. As shown in Haeringer and Hałaburda (2016), however, strategy-proofness itself is equivalent to this seemingly stronger requirement.

(2) If $\mathbb{S}^a\left(h^A\right) \neq \mathbb{S}^a\left(h^{A\prime}\right)$, then h^A and $h^{A\prime}$ must be in different elements of $\mathcal{I}^a$.

In other words, if two collective histories are such that the agent a's choice histories are different, or if they result in different menus to be given to a afterwards, these must be in different elements of the partition $\mathcal{I}^a$. The partition $\mathcal{I}^a$ represents agent a's information sets, and reflects the collective histories that an agent can differentiate based on what she can observe *and by the information that is provided by the specific implementation of the PAO mechanism*. Collective histories in the same element of the partition cannot be differentiated explicitly based on the information that the agent is given. The restriction that we defined above is the minimal information that we assume the agents must have: they know the menus that were offered to them and the choices they made. Our results will be robust to finer partitions, including that of perfect information, in which each partition element contains only one collective history. We denote by $\mathcal{I} = (\mathcal{I}^{a_1}, \ldots, \mathcal{I}^{a_n})$ the **information structure of the game** being considered. Given an information structure $\mathcal{I}$, a **belief system** θ is a collection of probability measures, one for each element of the partition in the information structures in $\mathcal{I}$. We denote by $\theta^a\left(h^A\right)$ the probability associated with collective history h^A in $\mathcal{I}^a$.

Next, we define an agent a's **strategy** to be a function $\sigma_a : \mathcal{I}^a \rightarrow O$, where for every $\mathcal{I}_i^a \in \mathcal{I}^a$, and $h^A \in \mathcal{I}_i^a$, $\sigma_a\left(h^A\right) \in \mathbb{S}^a\left(h^A\right)$. That is, strategies map from information sets to objects in the menu offered to the agent.[21] A collection with one strategy per agent is a strategy profile $\sigma = (\sigma_{a_1}, \ldots, \sigma_{a_n})$.

Fix a belief system θ, and let $a \in A$ be an agent and h^A be a collective history in $H_{\mathbb{S}}^A$. Let $\mathcal{I}_i^a$ be the set in agent a's information structure such that $h^A \in \mathcal{I}_i^a$. We define the **outcome belief for** a **under** θ, $\mathcal{O}_a^\theta\left(h^A, \sigma\right)$, as the distribution over assignments that result from following the PAO mechanism $\mathbb{S}$, starting from the collective history h^A, in which agents follow the strategies in σ, given the distribution that θ puts over the elements of the set $\mathcal{I}_i^a$.

Let A and B be two random assignments. We denote by $\succcurlyeq_a$ the first-order stochastic dominance relation under P_a. That is, $A \succcurlyeq_a B$ if for all $o \in O$, $Pr\left\{A\left(a\right) = o' | o' R_a o\right\} \geq Pr\left\{B\left(a\right) = o' | o' R_a o\right\}$. We can now define an ordinal perfect Bayesian equilibrium (OPBE).

Definition 4. A strategy profile σ together with a belief system θ is an **ordinal perfect Bayesian equilibrium (OPBE)** if for every $a \in A$, every $h^A \in H_{\mathbb{S}}^A$, and every strategy ${\sigma_a}'$ for agent a:

$$\mathcal{O}_a^\theta\left(h^A, (\sigma_a, \sigma_{-a})\right) \succcurlyeq_a \mathcal{O}_a^\theta\left(h^A, ({\sigma_a}', \sigma_{-a})\right)$$

Even though an OPBE is a refinement of the concept of a perfect Bayesian equilibrium,[22] it may still suffer from the fact that an equilibrium may be supported by artificially

[21]For simplicity, and without any consequence, we allow for the definition of strategies to make "choices" in collective histories in which an agent is not given a menu.

[22]If a strategy profile σ together with a belief system θ is an OPBE, then σ together with θ is a perfect Bayesian equilibrium for any utility functions profile that represents the ordinal preferences in P.

constructed belief systems. The following alternative equilibrium notion improves upon that.

Definition 5. A strategy profile σ is a **robust ordinal perfect Bayesian equilibrium** if for every belief system θ, σ is an OPBE.

We say that an allocation rule φ is **pick-an-object implementable** in some equilibrium notion if there exists a PAO mechanism $\mathbb{S}$ that sequentializes φ, in which straightforward strategies constitute an equilibrium in that notion.

Theorem 2. *A rule is pick-an-object implementable in robust ordinal perfect Bayesian equilibrium if and only if it is strategy-proof and satisfies monotonic discoverability.*

3.1. **Relation with Obviously Strategy-proof Mechanisms.** When introducing the concept of obvious strategy-proofness (OSP), Li (2017) defined an OSP mechanism as one that has an equilibrium in obviously dominant strategies. A rule φ is OSP-implementable if there is a game and an obviously dominant strategy for each type of player in that game, such that the outcome produced by this strategy for each type profile is what is determined by φ. Pycia and Troyan (2019) showed that, in an object allocation environment such as the one that we use, every OSP-implementable rule could be implemented via a *millipede game* with a *greedy strategy*. Millipede games are sequential games where, in each period some agent can either "pass" or "clinch" one of potentially multiple options in a menu, which corresponds to private allocations that they can guarantee (in our setup, therefore, that agent can clinch an object and leave with it). A greedy strategy consists of an agent choosing to "pass" as long as her most-preferred object can still be clinched from a menu in some continuation history, and clinching it whenever it is in a given menu.[23]

The first thing to note is that there is a direct relation between a greedy strategy in a millipede game and a straightforward strategy in a PAO mechanism. In a millipede game, every time an agent interacts with the mechanism she is given a menu of objects that she can pick and keep for good and, potentially, also one option to "pass." That is, as in PAO mechanisms, in millipede games an agent is always matched to the last object that she chose. However, in a millipede game, an agent can only choose an object once, and it remains her final allocation. In contrast, in a PAO mechanism, a chosen object can be interpreted as her "tentative allocation," and thus she can potentially choose multiple times. The authors show that the greedy strategy is obviously dominant in that game.

[23]While each option in a menu given to a player in a millipede game corresponds to a private allocation for that player, there might be multiple options in that menu associated with the same allocation. While every option associated with an allocation results in the same outcome for the player making that choice, different choices among these might result in different outcomes for other players. Our results will consider only situations in which that is not the case: different items in a menu correspond to different private allocations for the agent making the choice. Therefore, we will ignore the possibility that they do not.

As long as an agent can infer the set of clinchable objects in continuation histories, following the greedy strategy is very simple: the only object she chooses from a menu is her most preferred in the set of feasible objects. Otherwise, she simply passes.[24]

In a canonical PAO mechanism, on the other hand, menus always contain the entire set of feasible objects. Therefore, an agent following a straightforward strategy chooses the most-preferred object among those that are feasible, and will only have to make another choice once, and then only if that object is no longer feasible. Notice that, as opposed to millipede games, in these PAO mechanisms agents do not have to infer the set of feasible objects: all feasible objects are offered in the menus. Our next result relies on the common property of non-bossiness, defined below.

Definition 6. A rule φ is **non-bossy** if, for every $a \in A$, if $\varphi_a(P_a, P^{-a}) = \varphi_a(P'_a, P^{-a})$, then $\varphi(P_a, P^{-a}) = \varphi(P'_a, P^{-a})$.

If a rule is non-bossy and OSP implementable, we obtain the following relation with PAO mechanisms:

Theorem 3. *Every non-bossy OSP implementable rule is PAO implementable in weakly dominant strategies.*

Given the simplicity of these strategies, a question one may have is whether what drives a high proportion of behavior in line with an obviously dominant strategy is, in fact, the simplicity of the strategy itself, as opposed to more sophisticated arguments in terms of the kind of counterfactual reasoning that is eliminated in an OSP mechanism. If that is the case, a PAO mechanism could also lead agents to behave more in line with its theoretical prediction.

4. Experiments

In this section, we present a series of experiments designed to test the performance of PAO mechanisms for implementing different allocation rules, when compared to the traditional direct revelation mechanisms and OSP mechanisms, when possible. We chose two rules: top trading cycles (TTC) and serial dictatorship (SD). Ideally, we would have liked to compare implementing multiple allocation rules via PAO versus multiple alternative mechanisms. Our choice of using SD and TTC is driven mainly by the importance of both rules for practical applications. Another contender could be DA, but the comparison of the PAO implementation of DA versus its direct counterpart was made in two recent papers Bó and Hakimov (2020) and Klijn *et al.* (2019). Both led to a similar conclusion, reporting the better performance of the PAO mechanism, with respect to the truth-telling

[24] Pycia and Troyan (2019) question, however, the simplicity of inferring the set of clinchable objects in continuation histories, as it requires foresight from agents. They then introduce the concept of strong obvious strategy-proofness. In a strong OSP mechanism, the agents face a menu to choose from only once, and there is no passing option. The obviously dominant strategy in these games thus simply requires choosing the best object from the single menu that is offered.

rates and stability of the allocations, when compared to the direct DA mechanism. Additionally, our choice of allocation rules was driven by the possibility of OSP implementation of both rules (in the case of acyclic priorities for TTC). This allows us to compare PAO versus OSP implementations in a restricted setup.

4.1. **Mechanisms.** In this subsection, we describe the six mechanisms used in the experiments. For each allocation rule, namely TTC and SD, we use three different mechanisms: Direct[25], PAO, and OSP. For the PAO mechanisms, we used the canonical mechanism induced by these rules (see section 2.3). The mechanisms correspond to treatments in the experiments. Note that we describe the mechanisms in the same way they were described to the participants. We omit the description of actions when some or all objects are not acceptable for simplicity and to prevent confusion by the subjects. The reason is that all objects lead to a positive payoff for the subjects in the experiments. Moreover, in the experiments, subjects had to choose at least one object in all sequential mechanisms and list all objects in the rank-order lists in the direct mechanisms. Details about the information given to the participants in the experiment are detailed in section 4.2.

Direct TTC

Every participant submits her rank-order list of objects to a central authority. The following steps are executed by the central authority, without any further participation from the subjects.

- Step 1.1 All participants point to the object at the top of their submitted rank-ordered lists. Each object points to the participant with the highest priority at that object.
  - The mechanism looks for cycles. There is at least one cycle. All participants in the cycle are assigned the object they pointed to.
- Step 1.2 The priorities of the objects are updated to account for assigned participants. Submitted rank-ordered lists are updated to account for assigned objects. Steps 1–1.2 are repeated until all objects are assigned.

Direct TTC is strategy-proof, and Pareto efficient (Abdulkadiroğlu and Sönmez, 2003).

PAO TTC

All participants are asked to pick one object from a menu with all objects.

- Step 1.1. All participants point to the object they picked. All objects point to the participant with the highest priority at that object.
  - The mechanism looks for cycles. There is at least one cycle. All participants in the cycle receive the object they pointed to.
- Step 1.2 For the remaining participants, if their last picked object was already assigned, the participant is asked to choose a new object from a menu of remaining objects. Steps 1–1.2 are repeated until all the objects are assigned.

[25]We capitalize "Direct" when referring to the treatment name.

In PAO-TTC, every participant following the straightforward strategy is a robust OPBE (see Theorem 2). If all participants follow straightforward strategies, the allocation is Pareto efficient.

OSP TTC

Note that OSP TTC is only defined when the priorities of the objects are acyclic. We use the mechanism described in Troyan (2019).

- Step 1.0. The mechanism first tentatively assigns each object to the participant with the highest priority at that object (i.e., the participant *tentatively owns* the object).
- Step 1.1. One by one each participant who tentatively owns an object is asked whether one of the objects that she owns is her favorite object. There are two possible answers for each participant and object :
  - If "Yes," the corresponding participant receives the object. Go to Step 1.2.
  - If "No," she is asked about the next object among the ones she tentatively owns. If the participant answers "No" to all tentatively owned objects, the algorithm moves to the next participant who owns at least one object. If all participants who tentatively own at least one object say "No" to all owned objects, then each participant who tentatively owns at least one object is asked to pick one object among the objects she does not own.
    - ∗ All participants point to the object they picked. Each object points to its owner. The mechanism looks for cycles. There is at least one cycle. All participants in the cycle receive the object they picked.
- Step 1.2 The priorities of the objects are updated to account for the participants who left. Steps 1.1–1.2 are repeated until all objects are assigned.

In OSP TTC, the truthful strategy[26] is obviously dominant, leading to a Pareto efficient allocation.

Note that in mechanisms implementing the SD allocation rule, we use priority scores instead of using ordinal tables priorities. In the experiments, each subject knew her score, and knew that the higher the score, the higher her priority. Details are explained in the experimental design section.

Direct SD

Every participant submits her rank-order list of objects to a central authority. The following steps are executed by the central authority, without any further input from the subjects.

[26]More specifically, the truthful strategy consists of only saying "Yes" when the object is the most-preferred among the objects that are still available, and picking the most preferred object when asked to pick from objects that she does not own.

- Step 1. The participant with the highest priority score is assigned the top-ranked object on her list.
- Step 2. The participant with the second-highest priority score is assigned the top object on her list, among the remaining objects. ....
- Step N. The participant with the Nth-highest priority score is assigned whatever object remains.

Direct SD is Pareto efficient and strategy-proof.

PAO SD

- Step 1. All participants are asked to pick one object from a menu with all objects.
- Step 2. The participant with the highest priority score is assigned the object she picked. All other participants who chose this object are asked to pick a new object from a menu containing the remaining objects.
- Step 3. The participant with the second-highest priority score is assigned the last object she picked. All other participants with lower priority who chose this object are asked to pick a new object from the menu containing the remaining objects. ....
- Step N+1. The participant with the lowest priority score is assigned the last object she picked.

In PAO-SD, every player using the straightforward strategy is a robust OPBE (see Theorem 2). If all participants follow straightforward strategies, the allocation is Pareto efficient.

OSP SD

- Step 1. The participant with the highest priority is asked to pick an object and she is assigned the object she picked.
- Step 2. The participant with the second-highest priority is asked to pick an object among the remaining ones, and she is assigned the object she picked. ....
- Step N. The participant with the lowest priority is assigned to the last object that remains.

In OSP SD, straightforward strategies are strongly obviously dominant, leading to a Pareto efficient allocation (Pycia and Troyan, 2019).

4.2. **Experimental design.** In the experiment, there were eight objects and eight participants. In all treatments, participants received 22 euros if they were matched to their most-preferred object, 19 euros to their second most-preferred object, 16 euros to their third most-preferred object, and so on. Participants received 1 euro if they were matched to their least-preferred object.

Each session lasted for 21 rounds. At the end of the experiment, one round was randomly drawn to determine the participants' payoffs.

Each round represented a new market. The preferences used in each market were generated following the designed market idea of Chen and Sönmez (2006). For each market, agents' ordinal preferences were generated from cardinal utilities by adding, for each object, a common and an idiosyncratic value. The common values for each object were drawn from the uniform distribution with the range $[0, 40]$. The idiosyncratic values were drawn, for each object and each agent, from the uniform distribution with the range $[0, 20]$. The agent's utility from being matched to the object was the sum of both components. The resulting utilities were transformed into ordinal preferences. The procedure above ensures some correlation between participants' preferences. All objects had priorities over participants. The priorities were independently drawn from uniform distributions, in each round.[27]

We ran three treatments between-subjects: Direct, PAO, and OSP. Comparing Direct and PAO is our main focus. We additionally ran OSP, which is only available in SD and for a restricted set of environments in TTC – namely those with acyclic priorities – to disentangle the effect of the sequential pick-an-object setup from the effect of the simpler strategic setup of OSP mechanisms. To have a comparison of the treatments under a variety of parameters, including different allocation rules, we ran three environments within-subjects: TTC with cyclic priorities, TTC with acyclic priorities, and SD.

Note that, with respect to the TTC rule, acyclic priorities is the most applicable environment, and therefore the main focus of our analysis for that rule. TTC with acyclic priorities is a simpler decision environment as the acyclic priorities make sure that, at any time of the mechanism, only two participants could be at the top of the priorities of all objects.[28] TTC with acyclic priorities is implementable through an OSP mechanism, thus allowing us to compare it to a PAO implementation of that rule. SD is arguably the simplest allocation rule and also allows implementation through the OSP mechanism, which also coincides with the setup of sequentially making choices from a menu. Unlike the general PAO mechanism, it does not require the simultaneous play of participants at the first step.

For the first 14 rounds of the experiment, participants were matched using the TTC rule. During the first seven of these rounds, they faced markets with cyclic priorities in the Direct and PAO treatments. Given the importance of this comparison, and to prevent learning effects from other environments implemented within-subjects, we always run cyclic TTC in the first seven rounds. The same markets were used for Direct and PAO treatments. Because there is no OSP mechanism for TTC with cyclic priorities, we also generated acyclic priorities for the first seven rounds of the OSP treatment. For rounds eight to 14, TTC with acyclic priorities was used to match participants to objects. The same markets were used in all three between-subjects treatments.

[27]For rounds with acyclic priorities, priorities were drawn from the set of acyclic priorities. A new draw of acyclic priorities was generated for each round.

[28]TTC cycles, therefore, could involve only one or two agents.

Rounds	Direct	PAO	OSP
1-7	TTC cyc	TTC cyc	TTC acyc
8-14	TTC acyc	TTC acyc	TTC acyc
15-21	SD	SD	SD

Table 1. Summary of treatments

In the first 14 rounds, participants observed the full priority tables of all objects. The provision of priority tables is necessary, as it allows participants to see when they are acyclic in the OSP TTC. Participants knew only their own preferences, however, and not the preferences of other participants. They knew that other participants might have the same or different preferences. We chose this informational environment to simplify the processing of market information, as providing complete information would lead to a longer decision time every round.

Finally, the SD allocation rule was used for the last seven rounds (rounds 15 to 21). Before round 15, subjects were informed about the switch of the mechanism. In the last seven rounds, the participants were assigned a priority score instead of observing the priority table of objects. They knew that the higher the score, the higher the priority. Participants were also informed that the priority scores would be drawn for each subject and each round independently from a uniform distribution with the range $[1, 100]$ and each participant was informed of their own draw. This choice was made to ensure that even the participants with the lowest score have incentives to play a meaningful strategy. Otherwise, they would know that their choices were irrelevant for the allocation. As for preferences, just like in the first 14 rounds, participants only knew their own preferences and were informed that other participants might have different preferences.

Table 1 presents the summary of the experimental design. Each cell of the table represents the mechanism used. Note that we explained to each subject two mechanisms to compare treatments under more than one allocation rule. It was feasible, as all SD mechanisms are quite simple and straightforward to explain; thus we decided to run it within-subjects together with the TTC mechanisms. After each round, the participants learned the object they were matched to, but not the matches of the other participants.

The within-subject design is driven by practical considerations and the goal of comparing treatments under various environments. We did not randomize the order of environments intentionally because we aimed at having high statistical power for the main comparison of the paper—PAO versus Direct under general priorities of TTC, given the number of participants. This environment was always run first and thus represents a clean comparison. The comparison between treatments under other environments might be affected by subjects' experience in the previous rounds. If subjects are more likely to play truthfully in the next round, given the truthful play in the previous round using the same kind of mechanism, it is possible that the design biases the results "against" the PAO treatment and "in favor" of the OSP treatment, since players in the OSP treatment play OSP TTC (with acyclic priorities) for 14 rounds, while in the PAO treatment, there were 7 rounds with acyclic priorities, followed by 7 rounds with cyclic priorities.

The experiment was run at the experimental economics lab at the Technical University of Berlin, from March to May 2019. We recruited student subjects from our pool with the help of ORSEE (Greiner, 2015). The experiments were programmed in z-Tree (Fischbacher, 2007). Independent sessions were carried out for each of the three between-subjects treatments. Each session consisted of either 16 or 24 participants that were split into two or three groups of eight for the entire session. We use fixed groups in order to increase the number of independent observations and allow for maximum learning. As every round represents a new market and subjects play under incomplete information about the preferences of the other participants, subjects cannot identify the strategies or identities of the players from previous rounds. Thus, we are not concerned about repeated games caveats.

In total, 15 sessions with 296 subjects were conducted. Thus, we have 96 subjects and 12 independent observations in Direct and PAO treatments, and 104 subjects and 13 independent observations for the OSP treatment. On average, the experiment lasted 110 minutes, and the average earnings per subject were 27 euros, including a show-up fee of nine euros.

At the beginning of the experiment, printed instructions were given to the participants (see Appendix B). Participants were informed that the experiment was about the study of decision-making. The instructions were identical for all participants of each treatment, explaining the experimental setting in detail. First, the mechanism implementing the TTC rule was explained, with an example. The participants were told that this mechanism would be used to match them to objects for the first 14 rounds. Then, the mechanism implementing the SD rule was explained, also with an example, and participants were told that this mechanism would be used to match them to objects in the last seven rounds. After round 14, participants were reminded of the switch of the mechanism. They were invited to re-read instructions for the second mechanism. Clarifying questions were answered in private. Note that the switch between cyclic and acyclic priorities does not switch the mechanism, and this change was not emphasized to subjects. They could, however, infer the difference from the priority tables.

4.3. **Experimental Results.** The significance level of all our results is 5%, unless otherwise stated. For all tests, we use the p-values of the coefficient of the treatment dummy in regressions on the variable of interest. Standard errors are clustered at the level of matching groups, and the sample is restricted to the treatments that are of interest for the test. We use the $>$ sign in the results between treatments to communicate significantly higher, and the $=$ sign to communicate the absence of statistically significant difference between treatments.

4.3.1. *Truthful reporting.* The focus of the paper is the comparison of the different proportions of equilibrium behavior induced by different mechanisms for the same environment.

We consider the proportions of subjects following a dominant strategy of truthful reporting in the direct mechanisms, a straightforward strategy that constitutes a robust OPBE in the PAO mechanisms, and the obviously dominant strategy in the OSP mechanisms. To simplify the language of distinguishing between these strategies, we use the concept of *truthful strategy.* A participant follows a truthful strategy under a direct mechanism when she submits the truthful list of all eight objects.[29] In PAO mechanisms and OSP-SD, a participant following the truthful strategy is equivalent to her following the straightforward strategy. In OSP-TTC, participants following the truthful strategy must truthfully answer the "yes-no" questions about the most-preferred object among the ones for which the participant has the highest priority, and in case of all "no" answers, the participant must make a truthful choice of the favourite object among the other objects.

Result 1 (Behavior in line with the truthful strategy):

(1) Under the TTC rule with cyclic priorities, the comparison of average proportions of subjects behaving in line with the truthful strategy leads to the following results: PAO>Direct.[30]
(2) Under the TTC rule with acyclic priorities, the comparison of average proportions of subjects behaving in line with the truthful strategy leads to the following results: OSP>PAO>Direct.[31]
(3) Under the SD rule, the comparison of average proportions of subjects behaving in line with the truthful strategy leads to the following results: OSP>PAO=Direct.[32]

Support:

Figure 2 presents the proportions of truthful strategies played by participants by treatments and rounds.

First, under the TTC rule with cyclic priorities, the average proportion of truthful strategies under direct TTC is 20 percentage points lower than under PAO TTC. The difference is significant. The significance of the difference is robust to modifications of the definition of truthful strategy in the direct TTC. Note that, in the setup of cyclic priorities, no mechanism that implements the TTC rule is OSP (Li, 2017). We observe

[29]Note that there are typically multiple undominated strategies, given the information available to the subjects taking part in the experiments. We argue that the submission of the full truthful list (truncations are not allowed by design), however, is the simplest strategy among the undominated. Nevertheless, as a robustness check we consider alternative definitions of truthful strategies throughout the section.

[30]The result is robust to two changes of the definition of truthful strategy in the direct TTC. First, if instead of requiring the full truthful list we count as truthful all truthful submissions until the guaranteed object, the result remains the same. Second, if instead of requiring the full truthful list in the direct TTC we count as truthful all submissions with the truthful ranking of objects until the assigned object, the result also remains the same.

[31]If instead of requiring the full truthful list in the direct TTC we count as truthful all truthful submissions until the guaranteed object, the result remains. If instead of requiring the full truthful list in the direct TTC we count as truthful all submissions with the truthful ranking of objects until the assigned object, the difference between Direct and PAO becomes not significant.

[32]If instead of requiring the full truthful list in the direct SD we count as truthful all truthful submissions until the assigned object, the result remains. Note that under the informational conditions of the last seven rounds, there is no such thing as "guaranteed objects."

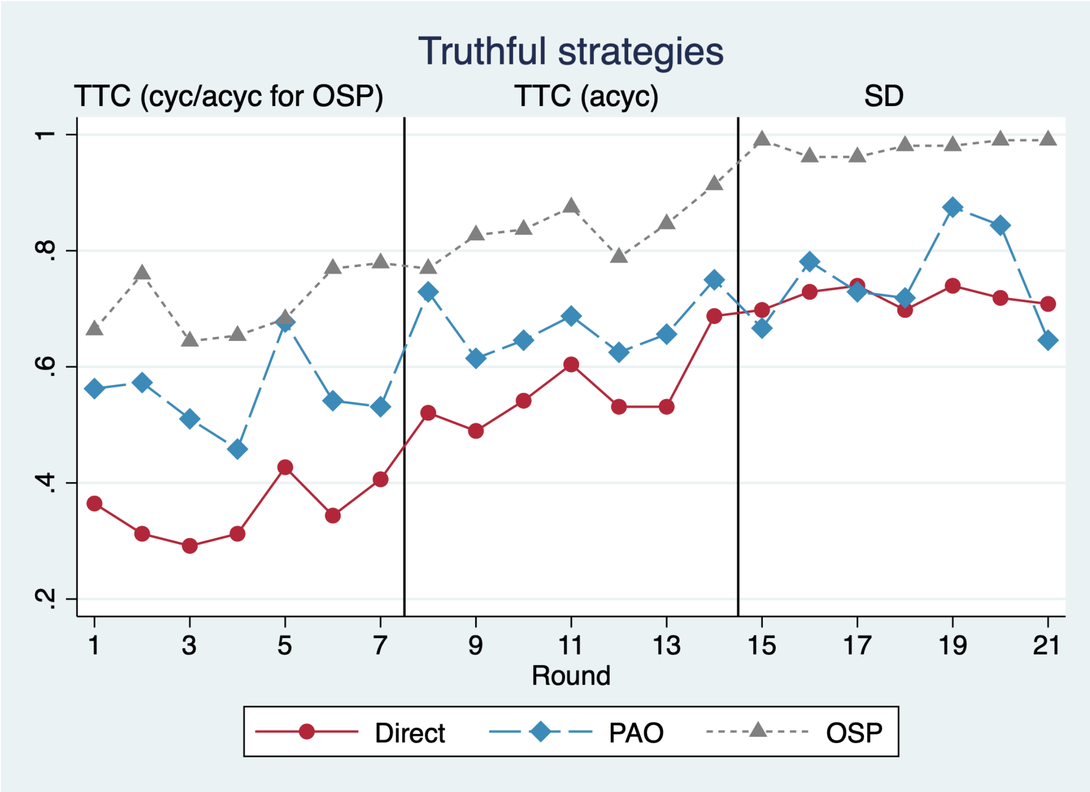

FIGURE 2. Truthful strategies by treatments and rounds

	Direct	PAO	OSP	Direct=PAO p-value	Direct=OSP p-value	PAO=OSP p-value
TTC cyclic	35%	55%	n/a	0.00	n/a	n/a
TTC acyclic	56%	67%	77%	0.00	0.00	0.00
SD	72%	75%	98%	0.47	0.00	0.00

Notes: All the p-values are for the coefficient of the dummy for the corresponding treatment in the probit regression of the dummy for the truthful strategy on the treatment dummy, with the sample restricted to the treatments involved in the test. The standard errors of all regressions are clustered at the level of the matching groups. Thus, we have 24 clusters in regressions comparing Direct and PAO treatments, and 25 clusters in regressions that involve a comparison of OSP treatment.

TABLE 2. Proportions of truthful strategies

that despite the strategy-proofness of direct TTC, the proportion of truthful strategies is only 35%, which is rather low. However, this rate is comparable to some other studies in the literature, that found similar rates of truthful reporting in TTC (for instance, 46% in Chen and Sönmez (2006) and 41% in Hakimov and Kesten (2018).)[33] Sequentialization of the mechanism through PAO leads to a significant increase in the proportion of truthful strategies. This finding is similar to the finding of Klijn *et al.* (2019) and Bó and Hakimov (2020), who show that the PAO implementation of the DA rule outperforms the direct mechanism. This is the main comparison of the experiment—in the absence of

[33]Two notable exceptions are Calsamiglia *et al.* (2010) and Pais and Pintér (2008), who documented higher rates of truth-telling under TTC (62% and 85%, respectively). The high rate in Pais and Pintér (2008) is likely driven by the fact that the rank-order list contained only three schools. Note, that in Hakimov and Kesten (2018) in an environment with five schools, the rate is just 30%.

OSP alternatives, the PAO implementation can improve the proportions of equilibrium behavior relative to the implementation through the direct mechanism.

Second, in the case of TTC with acyclic priorities, the OSP TTC outperforms both PAO and Direct mechanisms. The difference is significant for both the test considering only rounds 8 to 14, and the test considering rounds 1 to 14 in OSP TTC versus rounds 8 to 14 in Direct and PAO TTC.[34] As for the difference between Direct and PAO TTC, the proportion of truthful strategies is nine percentage points higher under PAO TTC, and the difference is significant.[35] The difference between Direct and PAO becomes smaller in TTC with acyclic priorities than in TTC with cyclic priorities. One can argue that learning is steeper under direct TTC. While this argument is hard to reject formally, suggestive evidence goes against this argument. More specifically, the coefficient for the variable "round" in the probit regression of the dummy for the truthful strategy is not significant in either Direct or PAO TTC. Also, there is a jump in the proportion of truthful strategies between rounds 7 and 8 under both treatments, suggesting that participants reacted to the switch of priorities from cyclic to acyclic.

Third, for the SD rule, the use of the OSP mechanism results in almost universal (98%) truthful behavior, as evident from Table 2. The rate of truthful strategies in OSP SD is significantly higher than in the Direct and PAO treatments. One can argue that the high rate of OSP is driven by the experience of subjects under OSP TTC, where the rate of truthful strategies was already higher. Again, the within-subjects feature of the design does not allow us rejecting this argument formally. However, note that the decision environment under SD was quite different due to the different informational environment, and we also notified subjects about the switch of the mechanism before round 15. Also, the rate of truthful reporting under OSP SD is quite similar to Hakimov *et al.* (2021b) who run OSP SD between-subjects. There is no significant difference between Direct and PAO. SD is a simple allocation rule, and thus the rates of manipulations are already relatively low under the Direct mechanism. Note that under PAO SD, participants might engage in multiple decisions, especially when they have a low priority: every time the chosen object is taken by someone with a higher priority, the agent is asked to pick another one.

Next, we take a closer look at the determinants of higher truthful rates under OSP.

Result 2 (The truthful strategy in OSP): In OSP TTC, the rate of truthful behavior in the passing actions is much lower than in the clinching actions, and the difference is significant.[36]

[34]One can argue that the difference is driven by learning, as subjects in the OSP treatment had already played the OSP mechanism in the first eight rounds. We acknowledge the bias of our design in favor of OSP treatment. Note, however, that the rate of truthful strategies in OSP treatment in the first seven rounds is higher than the rate of truthful strategies under Direct and PAO in rounds 14-20, suggesting that the difference is unlikely to be explained entirely by learning.

[35]Note, however, that the significance of the difference is not robust to redefining truthful strategies in direct TTC as the truthful ranking of objects in the rank-order lists until the assigned object. In this case, the difference between Direct and PAO becomes not significant.

[36]We follow the insights from Pycia and Troyan (2019) and define "passing actions" as OSP strategies, which involve saying "no" to all objects which the agent tentatively owns under OSP TTC. Thus, the

Support:

In line with the theory, obvious strategy-proofness leads to a higher rate of truthful strategies both under TTC and SD. Note that, in general, OSP-TTC is not strongly obviously strategy-proof, and the obviously dominant strategy might contain so-called "passing" actions that require forward-looking from participants. In line with the definition of Pycia and Troyan (2019), we can categorize possible paths of the obviously dominant strategies of the participants into two groups:

(1) **Clinching actions**—paths of the obviously dominant strategy that contain only clinching actions. If a participants was at the top of the priority of her favorite remaining object by the time the OSP TTC mechanism interviewed her (thus, she was at the top of the priority of at least one object). This strategy is strongly obviously strategy-proof, as it does not require passing to the other player, and thus does not require foresight from the participant.

(2) **Passing actions**—paths of the obviously dominant strategy that contain at least one passing action. If a participant was not at the top of the priority of her favorite remaining object by the time the OSP TTC mechanism interviewed her (thus, she was at the top of the priority of at least one object). OSP requires to "pass" the turn of the interview to another participant. This path of the obviously dominant strategy does require foresight from the participant. In the context of OSP TTC, these strategies require saying "no" to all objects and then picking the favorite object among those for which the participant does not have the top priority.

Table 3 presents the proportion of truthful strategies in OSP TTC, depending on the path of the OSP strategy. The rate of truthfulness in the passing actions is much lower than in the clinching actions. In fact, the rate of truthfulness in passing actions is not significantly different from truthful rates under direct TTC (p=0.94), and is significantly lower than under PAO ($p < 0.01$), which are not obviously strategy-proof. In contrast, once the path of OSP contains only clinching actions like in strongly obviously strategy-proof strategies, the truthful rate is much higher than in other treatments and reaches 93%.

	OSP TTC acyclic	
	N	% of truthful
Clinching actions	836	93%
Passing actions	620	56%

TABLE 3. Truthful strategies by clinching and passing in OSP TTC

This result supports the concept of strong obvious strategy-proofness by Pycia and Troyan (2019). Indeed, when the market is such that the preferences and priorities of the object are strongly negatively correlated (the agents prefer objects that rank them the

agent has to pass the decision to the next agent, hoping that a better object will appear in his choice set. "Clinching actions" are OSP strategies that require only saying "yes" to the most-preferred object in OSP TTC. Thus, whenever an agent is asked to act, she tentatively owns the most-preferred object.

lowest), the obvious strategy-proofness of OSP TTC might not result in the high rates of optimal behavior, as most paths of the OSP strategy will contain passing actions.

Summing up the subsection on individual strategies, experimental results support using the PAO mechanisms in the complex environment, where an OSP mechanism is not available. Once the environment is simple enough to allow for the presence of the OSP mechanisms, they should be used. The benefit of the OSP mechanisms comes mostly through the presence of the paths in OSP that contain only clinching actions.

4.3.2. *Efficiency.* While our experiment focuses on the subjects' individual behavior, we look at the efficiency of the reached allocations in this subsection. On one side, the matching game is often a zero-sum game, and a difference in efficiency is unlikely to appear.[37] Thus, as in many previous experiments, we do not expect to find large differences in efficiency between treatments. The differences might appear only due to Pareto-dominated allocations.

Result 3 (Efficiency):

(1) Under the TTC rule with cyclic priorities: Direct>PAO.
(2) Under the TTC rule with acyclic priorities: Direct>PAO, OSP = PAO (round 8-14), OSP = Direct (round 8-14).
(3) Under the SD rule : Direct>OSP=PAO.

Support:

Figure 3 shows the average rank of the assigned objects under the true preferences of participants by rounds. Thus, the higher the rank, the worse the assignment for participants. Table 4 shows the average ranks of assigned objects by treatments. Under TTC with cyclic priorities (first seven rounds), the average rank of the assigned objects is significantly higher under Direct than under PAO. The difference is, on average, 0.29 of a rank. This is a large difference for matching markets experiments, as often a worse assignment of one participant leads to a better assignment of the other. Note that we do not present the results for the OSP treatment for the first seven rounds. Because the comparison would not be meaningful, as the participants played under different priorities, and thus the equilibrium allocations are different.

Under TTC with acyclic priorities, there is a small but statistically significant difference between PAO and Direct. Again, the average assignment is significantly better for participants under the PAO treatment. There is no significant difference between OSP and other treatments. Thus, a higher rate of truthful strategies in OSP does not result in better average allocation of objects.

Finally, under the SD rule, the average rank in the Direct mechanism is significantly higher than under OSP and PAO. Despite a similar rate of truthful strategies, the consequence of the deviations from truthful strategies is different between Direct and PAO

[37]Consider the theoretical argument of (Liu and Pycia, 2016) for further details of this discussion.

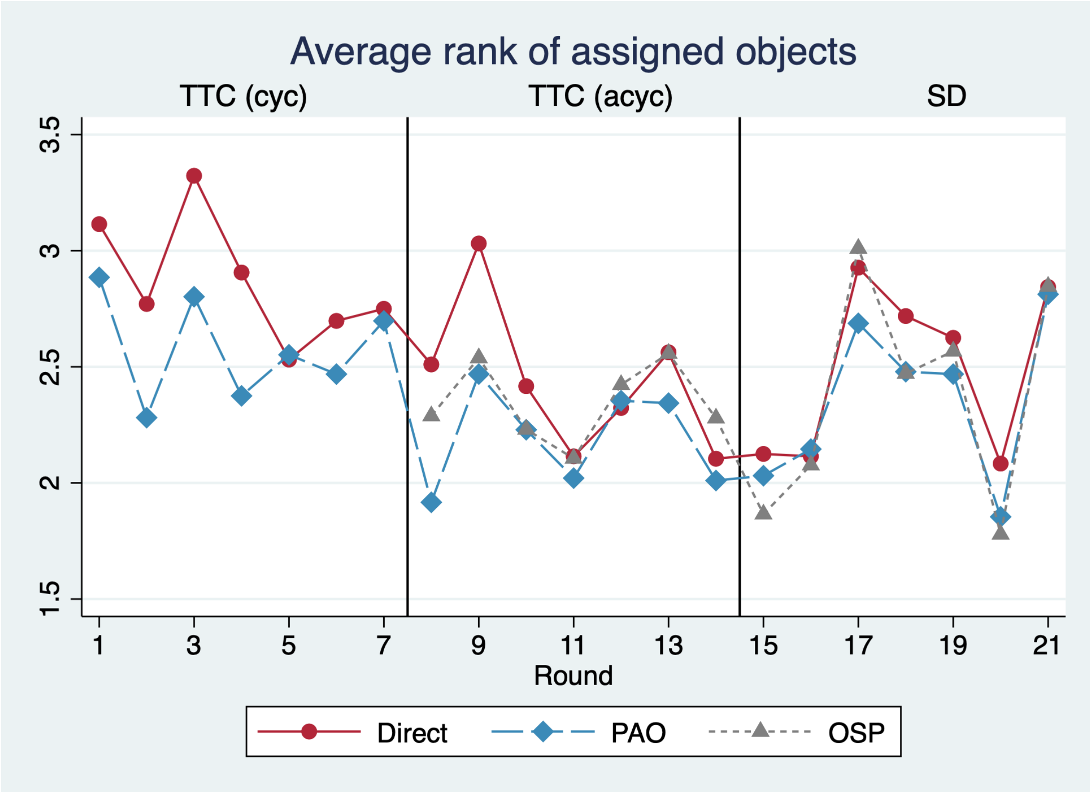

FIGURE 3. Average rank of assigned objects by treatments

	Direct	PAO	OSP	Direct=PAO p-value	Direct=OSP p-value	PAO=OSP p-value
TTC cyclic	2.87	2.58	n/a	0.00	n/a	n/a
TTC acyclic	2.43	2.19	2.34	0.00	0.33	0.08
SD	2.49	2.35	2.37	0.02	0.02	0.66

Notes: All the p-values are p-values for the coefficient of the dummy for the corresponding treatment in the OLS regression of the rank of the assigned object in the true preferences of participants on the treatment dummy, with the sample restricted to the treatments involved in the test. The standard errors of all regressions are clustered at the level of the matching groups. Thus, we have 24 clusters in regression comparing Direct and PAO treatments, and 25 clusters in regressions that involve comparison of OSP treatment.

TABLE 4. Average rank of assigned objects in the true preferences of the participants

treatments. At the same time, despite the large difference in truthful rates between PAO and OSP, the average rank of assigned objects does not differ significantly.[38]

Summing up the subsection on efficiency, PAO outperforms Direct in all environments. interestingly, PAO does not perform worse than OSP in both acyclyc TTC and SD, despite lower rate of truthful strategies. This is because in PAO some deviations from truthful strategies are payoff-irrelevant, while under OSP they are more likely to be payoff relevant under acyclyc TTC and are always payoff-relevant under SD.

[38]For alternative definitions of efficiency and estimation of the costs of deviation from truthful strategies in each treatment see Appendix B.0.1

5. Conclusion

Recent empirical evidence raises concerns about the practical success of strategy-proof matching mechanisms inducing truthful reporting of preferences in practical applications. Recent work by Li (2017) shed a new light on the design of market mechanisms, emphasizing the importance of simpler and thus potentially more successful solutions for practice. However, the hope was not long-lasting, as many desirable allocation rules cannot be implemented via obvious strategy-proof mechanisms.

Our paper takes a different stand on potential solutions to the perceived complexity of direct mechanisms. We suggest using PAO mechanisms when a rule cannot be implemented via OSP mechanisms but belongs to an extensive family of rules, which include many commonly considered in practice and the literature. Similarly to OSP mechanisms, PAO mechanisms can also implement those rules with an attractive and simple equilibrium strategy.

Our experimental evidence, together with recent evidence by Bó and Hakimov (2020) and Klijn *et al.* (2019), show that improvement over direct mechanisms in allocations and the percentage of people following an equilibrium is possible for allocation rules for which OSP implementation is not available. These results might appear puzzling, and may invite further research on understanding the relative strength of different equilibrium concepts in predicting behavior, especially of inexperienced participants.

References

Abdulkadiroğlu, A. and Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, **93** (3), 729–747. 1, 4.1

Akbarpour, M. and Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, **88** (2), 425–467. 1

Artemov, G., Che, Y.-K. and He, Y. (2017). Strategic 'mistakes': Implications for market design research*, Working Paper. 1

Ashlagi, I. and Gonczarowski, Y. A. (2018). Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, **177**, 405–425. 1

Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *American Economic Review*, **108** (2), 1452–1475. 1

— (2006). An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, **96** (3), 602–629. 1

Aygün, O. and Bó, I. (2021). College admission with multidimensional privileges: The brazilian affirmative action case. *American Economic Journal: Microeconomics*, **13** (3), 1–28. 1

Balinski, M. and Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, **84** (1), 73–94. 1

Bó, I. and Hakimov, R. (2016). The iterative deferred acceptance mechanism, Working Paper. 1, 7, 1, 19, 3

— and Hakimov, R. (2020). Iterative versus standard deferred acceptance: Experimental evidence. *The Economic Journal*, **130** (626), 356–392. 1, 6, 1, 4, 4.3.1, 5

Börgers, T. and Li, J. (2019). Strategically simple mechanisms. *Econometrica*, **87** (6), 2003–2035. 1

Breitmoser, Y. and Schweighofer-Kodritsch, S. (2021). Obviousness around the clock. 4, 1

Calsamiglia, C., Haeringer, G. and Klijn, F. (2010). Constrained school choice: An experimental study. *American Economic Review*, **100** (4), 1860–1874. 33

Chen, Y. and Sönmez, T. (2002). Improving efficiency of on-campus housing: An experimental study. *American Economic Review*, **92** (5), 1669–1686. 1

— and Sönmez, T. (2006). School choice: an experimental study. *Journal of Economic Theory*, **127** (1), 202–231. 1, 4.2, 4.3.1

Dur, U., Hammond, R. G. and Morrill, T. (2018). Identifying the harm of manipulable School-Choice mechanisms. *American Economic Journal: Economic Policy*, **10** (1), 187–213. 1

Echenique, F., Wilson, A. J. and Yariv, L. (2016). Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, **7** (2), 449–482. 1

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, **10** (2), 171–178. 4.2

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, **69** (1), 9–15. 2.2

Gong, B. and Liang, Y. (2016). A dynamic college admission mechanism in inner mongolia: Theory and experiment, Working Paper. 1, 7, 1

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, **1** (1), 114–125. 4.2

Grenet, J., He, Y. and Kübler, D. (2019). *Decentralizing centralized matching markets: Implications from early offers in university admissions*. Tech. rep., WZB Discussion Paper. 1, 1

Guillen, P. and Hakimov, R. (2017). Not quite the best response: truth-telling, strategy-proof matching, and the manipulation of others. *Experimental Economics*. 1

— and — (2018). The effectiveness of top-down advice in strategy-proof mechanisms: A field experiment. *European Economic Review*, **101**, 505–511. 1

Haeringer, G. and Hałaburda, H. (2016). Monotone strategyproofness. *Games and Economic Behavior*, **98**, 68–77. 20, A, 1

— and Iehle, V. (2019). Gradual college admission, Working Paper. 1

Hakimov, R., Heller, C., Kübler, D., Kurino, M. *et al.* (2021a). How to avoid black markets for appointments with online booking systems. *American Economic Review*, **111** (7), 2127–51. 1

— and Kesten, O. (2018). The equitable top trading cycles mechanism for school choice. *International Economic Review*, **59** (4), 2219–2258. 1, 4.3.1, 33

— and Kübler, D. (2021). Experiments on centralized school choice and college admissions: a survey. *Experimental Economics*, **24** (2), 434–488. 2

—, Kübler, D. and Pan, S. (2021b). Costly information acquisition in centralized matching markets. *Working paper*. 1, 4.3.1

— and Raghavan, M. (2020a). Transparency in centralised allocation. *Working paper*. 1

— and — (2020b). Transparency in centralised allocation, Working Paper. 1

Hassidim, A., Romm, A. and Shorrer, R. I. (2016). "Strategic" behavior in a strategy-proof environment. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, New York, NY, USA: ACM, pp. 763–764. 1

Jones, W. and Teytelboym, A. (2017). Matching systems for refugees. *Journal on Migration and Human Security*, **5** (3), 667–681. 1

Kagel, J. H., Harstad, R. M. and Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, **55** (6), 1275–1304. 1

Kawase, Y. and Bando, K. (2021). Subgame perfect equilibria under the deferred acceptance algorithm. 1

Klijn, F., Pais, J. and Vorsatz, M. (2019). Static versus dynamic deferred acceptance in school choice: Theory and experiment. *Games and Economic Behavior*, **113**, 147–163. 1, 1, 4, 4.3.1, 5

Li, S. (2017). Obviously Strategy-Proof mechanisms. *American Economic Review*, **107** (11), 3257–3287. 1, 3.1, 4.3.1, 5, B.0.1, B.0.1

Liu, Q. and Pycia, M. (2016). Ordinal efficiency, fairness, and incentives in large markets, Working Paper. 37

Luflade, M. (2018). The value of information in centralized school choice systems. *Working paper*. 1

Mackenzie, A. and Zhou, Y. (2020). Menu mechanisms, Working Paper. 1

Pais, J. and Pintér, Á. (2008). School choice and information: An experimental study on matching mechanisms. *Games and Economic Behavior*, **64** (1), 303–328. 1, 33

Pycia, M. and Troyan, P. (2019). A theory of simplicity in games and mechanism design, Working Paper. 1, 3.1, 24, 4.1, 36, 4.3.1, 4.3.1, A

Rees-Jones, A. (2018). Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *Games and Economic Behavior*, **108**, 317–330. 1

Roth, A. E. and Peranson, E. (1999). The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, **89** (4), 748–780. 1

—, Sönmez, T. and Ünver, M. U. (2004). Kidney exchange. *Quarterly Journal of Economics*, **119** (2), 457–488. 1

Schummer, J. and Velez, R. A. (2021). Sequential preference revelation in incomplete information settings. *American Economic Journal: Microeconomics*, **13** (1), 116–47. 1

Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, **1** (1), 23–37. 2.2

Shorrer, R. I. and Sóvágó, S. (2018). Obvious mistakes in a strategically simple college admissions environment: Causes and consequences, Working Paper. 1

Troyan, P. (2019). Obviously strategy-proof implementation of top trading cycles. *International Economic Review*, **60** (3), 1249–1261. 1, 4.1

Veski, A., Biró, P., Põder, K. and Lauri, T. (2017). Efficiency and fair access in kindergarten allocation policy design. *The Journal of Mechanism and Institution Design*, **2** (1), 57–104. 1

Appendix A. Proofs

Theorem 1. *There exists a pick-an-object mechanism that* ***sequentializes*** *an individually rational rule φ if and only if φ satisfies monotonic discoverability.*

Proof. First we assume that φ satisfies monotonic discoverability and show that the canonical pick-an-object function $\mathbb{S}$ is such that its mechanism **sequentializes** φ.

Define a pick-an-object function $\mathbb{S}$ such that for every i and $h^A \in H^A$:

$$\mathbb{S}^i\left(h^A\right) = \begin{cases} \emptyset & \text{if } \left|\varphi\left(h^A\right)\right| = 1 \text{ or } \overrightarrow{h_i^A} \in \mu_i^\varphi\left(h^A\right) \\ \mu_i^\varphi\left(h^A\right) & \text{otherwise} \end{cases}$$

To show that $\mathbb{S}$ sequentializes φ, we will show that: (i) if the collective history h^A is such that $\left|\varphi\left(h^A\right)\right| > 1$, at least one agent must be given a non-empty menu, that is, there must be at least one i such that $\mathbb{S}^i\left(h^A\right) \neq \emptyset$; (ii) When an empty menu is returned for all agents, their last choices must be the allocation that φ determines what should be produced by any preference profile consistent with the collective history.

First, (i). Suppose not. Then $\left|\varphi\left(h^A\right)\right| > 1$ and for all i, $\mathbb{S}^i\left(h^A\right) = \emptyset$. By the definition of $\mathbb{S}$, this implies that for all i, $\overrightarrow{h_i^A} \in \mu_i^\varphi\left(h^A\right)$. That is, the last choice of each agent is a feasible assignment after h^A.

Let $\overrightarrow{\mu}$ be the allocation that matches each agent with her last choice in h^A, that is, for every i, $\overrightarrow{\mu}(a_i) = \overrightarrow{h_i^A}$, and P^A be any preference profile consistent with h^A. Monotonic discoverability implies that either (a) for every preference profile P' consistent with continuations of h^A, $\varphi(P') = \overrightarrow{\mu}$, which is a contradiction with $\left|\varphi\left(h^A\right)\right| > 1$, or (b) that there is at least one agent a_{i*} such that for all $P \in \mathcal{L}\left(P^A, \overrightarrow{\mu}\right)$, $\overrightarrow{h_{i*}^A} \neq \mu_{i*}(P)$, which is again a contradiction with $\overrightarrow{h_{i*}^A} \in \mu_{i*}^\varphi\left(h^A\right)$.

Now, to (ii). Since an empty menu is given to all agents, then by the definition of $\mathbb{S}$, either (a) $\left|\varphi\left(h^A\right)\right| = 1$ or (b) for every i, $\overrightarrow{h_i^A} \in \mu_i^\varphi\left(h^A\right)$. Consider first (a). By definition of the notation, $\left|\varphi\left(h^A\right)\right| = 1$ implies that for all preference profiles consistent with h^A, the rule φ determines the same allocation. Suppose, however, that $\varphi\left(h^A\right) \neq \overrightarrow{\mu}$, and let a_i be an agent for whom $\varphi_i\left(h^A\right) \neq \overrightarrow{\mu}(a_i)$. Clearly, $\varphi_i\left(h^A\right)$ cannot be any choice made before $\overrightarrow{\mu}(a_i)$ in h^A, since by design of the pick-an-object mechanism it is only rejected if it is not a feasible assignment anymore. So it can either be an object which was present in a menu previous to the one where $\overrightarrow{h_i^A}$ was chosen but not in some future menu, or in the menu given when $\overrightarrow{h_i^A}$ was chosen. The first option contradicts the way in which menus are constructed: menus contain all feasible assignments conditional on the collective history. So if at some point the object type was not feasible anymore, it cannot be that agent's assignment under φ. For the second, this implies that the allocation was determined by φ to match agent a_i to some object type o^* that was not chosen. Notice, however, that since φ is individually rational, a collective history cannot point to a single allocation unless the agents' preferences consider these objects acceptable. That is, the collective history must include agents choosing, at some point, these objects from a menu that includes the option "$\emptyset$". Therefore, it cannot be the case that φ is individually rational, $\left|\varphi\left(h^A\right)\right| = 1$, and an $\varphi\left(h^A\right)$ matches some agent to an object that she did not choose from a menu. Now, case (b). Here, all the last choices of all agents are feasible in $\varphi\left(h^A\right)$. That is, there is no agent who will not be matched to their last choice for a preference profile that is consistent with $\varphi\left(h^A\right)$. But then monotonic discoverability implies that $\varphi\left(h^A\right) = \overrightarrow{\mu}$, which is what we wanted to show.

Next, we will show that if a rule φ does not satisfy monotonic discoverability, then it cannot be sequentialized by some pick-an-object function $\mathbb{S}$. Suppose not. Then there is a rule φ^* that does not satisfy monotonic discoverability and a pick-an-object mechanism $\mathbb{S}$ that sequentializes it. Since φ^* does not satisfy monotonic discoverability, there exists a preference profile P^* and an allocation μ^* such that (i*) $\varphi^*(P^*) \neq \mu^*$ and (ii*) for each agent $a \in A$, there is at least one preference profile $P^{*,a} \in \mathcal{L}(P^*, \mu^*)$ where $\varphi^*_a(P^{*,a}) = \mu^*(a)$. If there is more than one such allocation μ^* for that given P^*, let μ^* be such that for every $\mu' \neq \mu^*$ satisfying (i*) and (ii*), there is at least one $a \in A$ such that $\mu^*(a) P^*_a \mu'(a)$.

The first thing to note next is that (ii*) implies that when all agents follow straightforward strategies with respect to P^*, there must be for each agent a a period in which she chooses $\mu^*(a)$, and that this must happen before she chooses other objects that she is matched to by φ^* for any profile in $\mathcal{L}(P^*, \mu^*)$. That is, since when following P^* "up to μ^*" there is some continuation in which a is matched to $\mu^*(a)$, then a must first choose that object from a given menu.

The next observation is that, after the period in which a chooses $\mu^*(a)$, the determination of whether an agent a will be matched to $\mu^*(a)$ or some other object below $\mu^*(a)$ in her preference cannot depend on that agent's preferences among objects below $\mu^*(a)$ in her preference, since in order to obtain information about that part of agent a's preferences *requires rejecting* $\mu^*(a)$ *as a potential allocation for* a. This implies that whether a will be matched to $\mu^*(a)$ or not depends on information about the *other agents*' preferences. More than that, (ii*) specifically implies that the conclusion that a will not be matched to $\mu^*(a)$ cannot be reached before some other agent b makes choices *after* having her choice of $\mu^*(b)$ rejected. This, therefore, has the following implications:

- Every agent a^* following the straightforward strategy with respect to P^* will, in some period, choose her allocation under μ^*,
- In any periods that follow, in which the other agents did not yet choose their allocation under μ^*, agent a^*'s allocation may or not be determined to be $\mu^*(a^*)$, but will **not** have her choice rejected, since there is still some continuation in which $\mu^*(a^*)$ will be her allocation, and rejections are final.

The two implications above result in the following dynamic when agents follow straightforward strategies with respect to P^*: agents make choices over menus until, at some point, they choose their allocation under μ^*. After a certain number of periods, therefore, we reach a point in which all agents' last choices are their allocations in μ^*. By (i*), μ^*

should not be the allocation, but by (ii*), for each agent, more information about the preference profile is necessary to point out the correct allocation to be produced. This requires rejecting at least one of the agents' choices, but by (ii) for every agent there is a continuation in which she is matched to her assignment under μ^*. So no more information can be obtained when using any pick-an-object function, leading to a contradiction.

□

Proposition 1. *If φ is described by a generalized DA procedure, then φ satisfies monotonic discoverability.*

Proof. In light of Theorem 1, it suffices to show that there is a pick-an-object mechanism that sequentializes the rule that is specified by the generalized DA procedure. Let Ψ^* be the update function used to describe the rule φ. We construct the menu function $\mathbb{S}^*$ as follows:

$$\mathbb{S}^*\left(h^{A-\emptyset}\right) = (O, O, O, \ldots)$$

The value of $\mathbb{S}^*$ for other collective histories are determined, recursively, as follows. Let h^A be a collective history for which the value of $\mathbb{S}^*\left(h^A\right)$ has already been determined as $\mathbb{S}^*\left(h^A\right) = (\phi_1, \phi_2, \ldots, \phi_n)$.

For each choice profile $(o^1, o^2, \ldots, o^n)$, where for each $i \in A$, $o^i \in \phi_i$ if $\phi_i \neq \emptyset$ and $o^i = \diamondsuit$ otherwise,[39] perform the following:

(1) Construct the assignments μ^1 and μ^2 as follows:
  - For every $a_i \in A$:
    - If $o^i = \diamondsuit$ and $h^A \neq h^{A-\emptyset}$, let $\mu^1(a_i) = \overrightarrow{h_i^A}$.
    - Otherwise, let $\mu^1(a_i) = \emptyset$.
  - For every $a_i \in A$, let $\mu^2(a_i) = \emptyset$ if $o^i = \diamondsuit$, and $\mu^2(a_i) = o^i$ otherwise.

(2) For every $a_i \in A$, define the choice history h_i to be:
  - $h_i = h_i^A \oplus (\phi_i, o^i)$ if $o^i \neq \diamondsuit$,
  - $h_i = h_i^A$ otherwise.

(3) Let $\mu^3 = \Psi^*(\mu^1, \mu^2)$.
  - If for every $a \in A$ it is the case that $\mu^t(3) \in \{\mu^1(a), \mu^2(a)\}$, then $\mathbb{S}^*(h_1, h_2, \ldots, h_n) = (\emptyset, \emptyset, \ldots, \emptyset)$.
  - Otherwise, $\mathbb{S}^*(h_1, h_2, \ldots, h_n) = (\phi_1', \phi_2', \ldots, \phi_n')$, where for each $a_i \in A$:
    - $\phi_i' = \emptyset$ if $\mu^3(a_i) \in \{\mu^1(a_i), \mu^2(a_i)\}$,
    - $\phi_i' = O \setminus \bigcup_{(\Omega,\omega)\in h_i^A} \omega$ otherwise.

Notice first that $\mathbb{S}^*$ is a menu function: initial menus are the entire set O, and for every collective history the menus given are precisely the last menu an agent was given with her last choice removed. Since $\mathbb{S}^*$ is defined recursively for each collective history that can be generated by choices from the menus that can be offered, every possible path of the

[39] Here the symbol $\diamondsuit$ is used as a placeholder for the agents who are not presented with a menu to choose from.

pick-an-object mechanism is well defined. Consider now the pick-an-object mechanism $\mathbb{S}^*$ and the agents in A following straightforward strategies with respect to P.

What follows is an exact reproduction of the steps of the generalized DA procedure under the preference profile P. Since the menus always include all elements of O minus the agents' past choices, agents will make choices from the top until their last choice following their preference, as in the generalized DA. Agents who have their last choices rejected and are given new menus, for any collective history, are determined by the function Ψ^*, so that the way in which the sequences of choices from menus determine whether an agent is tentatively matched or not is given by that function. And finally, whenever the last choice of agents should be determined as the outcome, $\mathbb{S}^*$ returns a list of empty menus.

We can conclude, therefore, that the pick-an-object mechanism $\mathbb{S}^*$ sequentializes the rule φ. □

Theorem 2. *A rule is pick-an-object implementable in a robust ordinal perfect Bayesian equilibrium if and only if it is strategy-proof and satisfies monotonic discoverability.*

Proof. Let φ be a rule that is strategy-proof and satisfies monotonic discoverability. By Theorem 1, the mechanism $\mathbb{S}^\varphi$, described in the proof of that theorem, sequentializes φ.

That is, given some preference profile P, if every agent follows the straightforward strategy, the pick-an-object mechanism $\mathbb{S}^\varphi$ produces the outcome $\varphi(P)$.

For any set $I \subseteq O$, let $I!$ denote the set of all permutations of the elements of I. For any tuple γ of distinct elements of O, let:

$$\mathbb{P}|_\gamma \equiv \bigcup_{\lambda \in (O \setminus \gamma)!} \gamma \oplus \lambda$$

That is, $\mathbb{P}|_\gamma$ is the set of all preferences in which γ are the most-preferred object types, ordered as in the tuple γ itself. Let also $P|_I$, where $I \subseteq O$, be the preference P restricted to I.

We will use the following claim:

Claim. *Let $\mathbb{S}^\varphi$ be a pick-an-object mechanism, and h^A be any collective history in $H^A_{\mathbb{S}^\varphi}$. Then, there is a list of tuples $(\gamma_1, \gamma_2, \ldots, \gamma_n)$, each with distinct elements of O, such that the set of outcomes produced by $\mathbb{S}^\varphi$ in any continuation history of h^A is given by:*

$$\bigcup_{P \in \mathbb{P}|_{\gamma_1} \times \cdots \times \mathbb{P}|_{\gamma_n}} \varphi(P)$$

Proof. Given the definition of $\mathbb{S}^\varphi$, agents' choices are used as revealed preference, and the allocation that is produced is one that φ indicates being the unique allocation for the preference profiles consistent with these choices. By assumption, all agents followed some arbitrary strategy up to the collective history h^A.

When considering any agent $a \in A$, after the menus that were given to her and her choices, present in collective history h^A, there are potentially multiple preferences over the objects in $O \backslash \mu_a^{\varphi}\left(h^A\right)$ that are consistent with the choices made by a. One thing we can say, however, is that none of the objects in $\mu_a^{\varphi}\left(h^A\right)$ were chosen, since all of them may still become an outcome for a in a continuation collective history, by the definition of μ_a^{φ} itself. We can partition O into three sets O_1^a, O_2^a, O_3^a: O_1^a being the set of objects that a chose from a menu in some period in h^A, $O_2^a = \mu_a^{\varphi}\left(h^A\right)$, and O_3^a having all the other objects.

Clearly, any preference P_a in which:

- (i) $o \in O_1^a$ and $o' \in O_2^a$ implies $oP_a o'$ and
- (ii) $o \in O_1^a$, $o' \in O_1^a$, and o was chosen by a in h^A before o' implies $oP_a o'$

is consistent with the collective history h^A. Moreover, *any* continuation history of h^A is consistent with preferences with the following structure:

$$o_{1,1}^a \; P_a \; o_{1,2}^a \; P_a \cdots P_a \; o_{1,k}^a \; P_a \; \{O_3^a\} \; P_a \; \{O_2^a\}$$

where $\left(o_{1,1}^a, o_{1,2}^a, \ldots o_{1,k}^a\right)$ is the set O_1^a ordered by the period in which the object was chosen by a from a menu in , and $\{O_3^a\}$ and $\{O_2^a\}$ represent any permutation of the elements in these sets.[40] Since, for each agent a, all menus and choices that will take place in any continuation collective history involve only objects in O_2^a, continuation histories will be consistent with any particular ordering of the objects in O_3^a.

For every a and O_3^a defined above, let $o_{3,1}^a, o_{3,2}^a, \ldots, o_{3,\ell}^a$ be a permutation of the elements of O_3^a. The reasoning above allows us to conclude, therefore, that conditional on h^A, the outcome that will be produced by the pick-an-object mechanisms in question is $\varphi\left(P^*\right)$, where for each agent $a \in A$ its preference in the profile P^* is the following:

$$P_a^* = \gamma_a \oplus \lambda_a$$

where:

$$\gamma_a = \left(o_{1,1}^a, o_{1,2}^a, \ldots, o_{1,k}^a, o_{3,1}^a, o_{3,2}^a, \ldots, o_{3,\ell}^a\right)$$

and λ_a is a permutation of the elements of O_2^a consistent with the choices made by a, after the collective history h^A, over subsets of O_2^a. Since $\mathbb{S}^{\varphi}$ sequentializes φ, any combination of values of $\left(\lambda_a\right)_{a\in A}$ above which lead to different outcomes under φ are continuation collective histories of h^A in $H_{\mathbb{S}^{\varphi}}^A$. Therefore, for every $P \in \mathbb{P}|_{\gamma_1} \times \cdots \times \mathbb{P}|_{\gamma_n}$ and $\mu = \varphi(P)$, there is a collection of tuples $\left(\lambda_a\right)_{a\in A}$ such that the preference profile P^* constructed above is such that $\mu = \varphi(P) = \varphi\left(P^*\right)$, finishing the proof. □

We will also use the following result, which is a corollary from Theorem 1 in Haeringer and Hałaburda (2016):

[40]Note that preferences in which the objects in O_3^a are intertwined between the objects in O_1^a and O_2^a are also consistent with h^A. But since the pick-an-object mechanism being considered produces an outcome without eliciting more information, we can safely conclude that the allocation rule yields the same allocation for any preference consistent with those choices, in particular those preferences in which all objects in O_1^a are preferred to all objects in O_3^a, which in turn are all preferred to the objects in O_2^a.

Corollary 1. *[Haeringer and Haɫaburda (2016)] For every agent a, preference P_a, preferences of other agents P^{-a}, set $I \subseteq O$, permutation γ of I and $P^* \in \mathbb{P}|_\gamma$:*

$$\varphi_a \left(\gamma \oplus P|_{O\setminus I}, P^{-a}\right) R_a \varphi_a \left(P^*, P^{-a}\right)$$

Suppose now, for contradiction, that φ is not implementable in robust OPBE. Then, there is a belief system θ for which a strategy profile in which all agents follow straightforward strategies is not an OPBE. That is, there is an agent a and a strategy σ'_a, which is not straightforward after h^A, for which $\mathcal{O}^\theta_a \left(h^A, (\sigma_a, \sigma_{-a})\right)$ does not first-order stochastically dominate $\mathcal{O}^\theta_a \left(h^A, (\sigma'_a, \sigma_{-a})\right)$, where (σ_a, σ_{-a}) is the strategy profile in which all agents follow strategies that are straightforward in any continuation of h^A. That is, if $A = \mathcal{O}^\theta_a \left(h^A, (\sigma_a, \sigma_{-a})\right)$ and $A' = \mathcal{O}^\theta_a \left(h^A, (\sigma'_a, \sigma_{-a})\right)$, there is an object $o \in O$ for which:

$$Pr\left\{A = o' | o' R_a o\right\} < Pr\left\{A' = o' | o' R_a o\right\}$$

That is, the probability of obtaining an object at least as good as o is strictly higher under A' than under A. Since these random outcomes are produced by a distribution over deterministic outcomes, this in turn implies that there is at least one collective history $h^{A*} \in \mathcal{I}^a_i$, where $h^A \in \mathcal{I}^a_i$[41] and $\theta^a \left(h^{A*}\right) > 0$, in which the deterministic outcome of following the strategy profile (σ'_a, σ_{-a}) after h^{A*} is strictly preferred by a over the deterministic outcome that a obtains under the profile (σ_a, σ_{-a}) after h^{A*}. Let θ^* be a belief system in which $\theta^a \left(h^{A*}\right) = 1$, $o = \mathcal{O}^{\theta^*}_a \left(h^{A*}, (\sigma_a, \sigma_{-a})\right)$ and $o' = \mathcal{O}^{\theta^*}_a \left(h^{A*}, (\sigma'_a, \sigma_{-a})\right)$ (notice that since θ^* is degenerate at the information set $\mathcal{I}^a_i$, outcomes are deterministic for any given strategy profile). Since $o' P_a o$, the outcome of following the strategy profile (σ'_a, σ_{-a}) is different from following (σ_a, σ_{-a}), implying that under the first profile, agent a makes at least one choice from a menu that is not straightforward. That is, agent a chose a less preferred object with respect to P_a than another that was in the menu.

As by the claim above, the outcome of following the profile (σ_a, σ_{-a}) after h^{A*} is $\varphi\left(P^*_a, P^*_{-a}\right)$, whereas following the profile (σ'_a, σ_{-a}) after h^{A*} yields $\varphi\left(P^{**}_a, P^*_{-a}\right)$, where P^*_a and P^{**}_a differ only in how they rank the objects in $O\setminus\gamma_a$, all of them being at the tail of a preference ranking that is the same for all the remaining objects. Since P^*_a ranks the objects in $O\setminus\gamma_a$ with respect to P_a and P^{**}_a does not, agent a obtaining a more preferred object under $\varphi\left(P^{**}_a, P^*_{-a}\right)$ contradicts the recursive dominance of φ.

Finally, suppose that there is a rule φ^* which is *not* strategy-proof but is pick-an-object implementable in a robust OPBE. Let the pick-an-object mechanism $\mathbb{S}^{\varphi^*}$ implement φ^* in a robust OPBE. By theorem 1, φ^* satisfies monotonic discoverability. Since φ^* is not strategy-proof, then there is a preference profile P, an agent $a \in A$ and a preference $P'_a \neq P_a$ for which:

$$\varphi_a\left(P'_a, P_{-a}\right)\ P_a\ \varphi_a\left(P_a, P_{-a}\right)$$

[41]That is, h^{A*} and h^A are in the same information set.

For any agent a, let σ_a denote the straightforward strategy with respect to P_a, and σ the corresponding strategy profile. Since $\mathbb{S}^{\varphi^*}$ sequentializes φ^*, when all agents follow σ, the outcome produced by $\mathbb{S}^{\varphi^*}$ is $\varphi_a(P_a, P_{-a})$. Moreover, the outcome $\varphi_a(P'_a, P_{-a})$ is produced when all agents but a follow straightforward strategies, and a follows the deviation strategy σ'_a, "straightforward as if her preference was P'_a". Clearly, since the outcome is different, when following the deviation strategy σ'_a, agent a makes at least one choice that is not straightforward. Consider next the earliest collective history h^A in which the choices under the profiles σ and (σ'_a, σ_{-a}) differ, and a belief system θ for which $\theta^a\left(h^A\right) = 1$. By assumption, by following σ'_a instead of σ_a after h^A, agent a is strictly better off. Which is a contradiction with σ_a first-order stochastically dominating any other strategy, including σ'_a.[42] □

Theorem 3. *Every non-bossy OSP implementable rule is pick-an-object implementable in weakly dominant strategies.*

Proof. We prove the result by constructing a pick-an-object mechanism which serves as an interface between the agents and the millipede game of the OSP mechanism. Let φ^{OSP} be an OSP-implementable rule. Then, by Pycia and Troyan (2019), there is a millipede game Γ where, for every preference profile P, each agent following a greedy strategy with respect to their preference, results in the allocation $\varphi^{OSP}(P)$. In the game Γ, Nature moves once, at history $h^{\emptyset}$[43], and all players have perfect information. For every non-terminal history h (except for the one in which Nature moves,) there is an associated player $P(h) \in A$ and a menu $\phi(h)$.

By the definition of the millipede game in Pycia and Troyan (2019), the menu $\phi(h)$ might contain multiple items, each of them associated with an element of O, and at most one *pass* option. In principle, there might be multiple items associates with the same $o \in O$, and the choice among these options be consequential to the final allocation for other players. The characterization of OSP mechanisms as milipede games associated with a greedy strategy allows for that selection to depend on the preferences of the player making that choice. For example, you could have an agent a who receives a menu $\phi(h)$ containing, among other options, items ω and ω'. Choosing either will result in a being assigned object o, but the greedy strategy associated with this millipede game indicates that the agent chooses ω if her preference is $oP_ao'P_ao''$, but ω' if her preference is $oP_ao''P_ao'$.

Non-bossiness, however, implies that even if Γ contains multiple items for the same private allocation that are chosen, under the greedy strategy, when the agent has different preferences, these must result in the same allocation, fixed the other agents' greedy strategies. Therefore, every item associated with a private allocation is equivalent to

[42]Notice that since outcomes are deterministic, first-order stochastic domination here translates into no deviation leading to better outcomes than σ_a/

[43]If Nature does not move, then we can simply consider that Nature chooses a constant action at $h^{\emptyset}$.

each other. This allows us to assume, instead that items in the menu given to the players in a millipede game consist are all objects, with at most one pass option. That is, $\phi(h) \subseteq O \cup \{pass\}$.

With that, we can define the set of clinchable objects for agent a at continuations of h as follows:

$$C_a(h) = \bigcup_{h' \in P^{-1}(a): h \subseteq h'} \phi(h') \backslash \{pass\}$$

That is, $C_a(h)$ is the set of objects that are given in a menu to player a at some continuation history from h. The menu function $\mathbb{S}$ for the pick-an-object mechanism that we are constructing is such that $\mathbb{S}\left(h^{A-\emptyset}\right) = \left(C_{a_1}\left(h^{\emptyset}\right), C_{a_2}\left(h^{\emptyset}\right), \ldots, C_{a_n}\left(h^{\emptyset}\right)\right)$. The value of $\mathbb{S}\left(h^{A}\right)$ for any collective history $h^A = (h_1, h_2, \ldots, h_n)$ is defined as follows:

- Let h_i^j be the jth item in h_i, and $h_i^j = \left(\Omega_j^i, \omega_j^i\right)$. That is, Ω_j^i is the jth menu faced by agent a_i, in which she chose the item $\omega_j^i \in \Omega_j^i$.
- Let $h = h^{\emptyset}$ and for all $i = 1, \ldots, n$, let $\eta_i = 1$. Follow the game tree in Γ by making the agents play as follows:
  - **Step 0**: If h is a terminal node of Γ, then we can determine the value of $\mathbb{S}\left(h^A\right)$ to be $(\emptyset, \emptyset, \ldots, \emptyset)$.
  - **Step 1**: Let $a_i = P(h)$. That is, a_i is the active player at history h. If $\omega_{\eta_i}^i \notin C_{a_i}(h)$ and $\eta_i < |h_i|$, increase the value of η_i by 1. Otherwise, if $\omega_{\eta_i}^i \notin C_{a_i}(h)$ and $\eta_i = |h_i|$ we can determine the value of $\mathbb{S}\left(h^A\right)$ to be $(\phi_1, \phi_2, \ldots, \phi_n)$, where $\phi_j = \emptyset$ if $\omega_{\eta_j}^j \in C_{a_j}(h)$, and $\phi_j = C_{a_j}(h)$ otherwise.
  - **Step 2**: If $\omega_{\eta_i}^i \in \phi(h)$, follow the node that represents choosing $\omega_{\eta_i}^i$, and let h be the history that follows that choice. Otherwise, choose the node $Pass$, and let h be the history that follows that choice. In either case, go back to step 0.

Notice that the procedure above will produce a list of menus for any collective history, even those which could never take place under a pick-an-object mechanism. To get our result, however, it suffices for us to show that collective histories that are generated by agents following straightforward strategies while interacting with the pick-an-object mechanism $\mathbb{S}$ described above are translated to greedy strategies from these same agents playing the game Γ.

Suppose then, for contradiction, that agents follow straightforward strategies while interacting with the pick-an-object mechanism $\mathbb{S}$ but that is not translated, in the description above, into these agents following greedy strategies. That must imply, therefore, that there is an agent $a^* \in A$ who in the procedure described above either (i) chooses $Pass$ at some history h^*, and the most-preferred element in $C_{a^*}(h^*)$ with respect to P_{a^*} is in $\phi(h^*)$, or (ii) chooses an object $o \in \phi(h^*)$ while there is another object $o' \in C_{a^*}(h^*)$ for which $o' P_{a^*} o$.

Fist, consider case (i). Notice that, by the definition of the value of $\mathbb{S}\left(h^A\right)$ and by the description of Step 1, every time an agent receives a menu of objects, that menu contains all objects in the continuation outcomes of the history that is reached by following the collective history that precedes the offering of that menu. Moreover, Step 1 also implies that whenever a history h in which the last choice made by a^* from a menu is not in $C_{a^*}(h)$, the procedure does not follow the game Γ after h. Therefore, at every history h, the last choice made by a^* in her choice history is her most-preferred object in $C_{a^*}(h)$ with respect to P_{a^*}. Step 2 requires that *Pass* would only be chosen if that last choice was **not** in $\phi(h^*)$, a contradiction.

Next, case (ii). In the previous paragraph we established that every time an agent receives a menu of objects, it contains all objects in the continuation outcomes. Since the agent chooses o in the game Γ, Step 2 and the way in which the pick-an-object mechanism works requires that o must have been the last choice made by a^*. Since both o and o' are in $C_{a^*}(h^*)$, they both must be in the menu from which a^* chose o. But that means that $o' P_{a^*} o$ and a^* chose o from a menu containing o', which is a contradiction with a^* following the straightforward strategy.

Finally, note that any deviation from straightforward strategies in the pick-an-object mechanism $\mathbb{S}$ either leads to the same paths through Γ as the greedy strategy or to different ones. Since the greedy strategy is obviously dominant in Γ, then it weakly dominates any other strategy, including any deviation induced by deviations from straightforward strategies in $\mathbb{S}$. Therefore, straightforward strategies are weakly dominant in the pick-an-object mechanism $\mathbb{S}$.[44] □

Appendix B. Additional experimental results

B.0.1. *Alternative efficiency definitions.* Given the fact that it might be the case that the equilibrium allocation has a lower sum of ranks than the allocation when a participant deviates from the truthful strategy, it is essential to look at different criteria. One measure of the success of the mechanism is whether the desired allocation is reached. This approach is used in Li (2017) to estimate the performance of direct versus OSP SD. We adopt this approach for all treatments. Note, however, that in Li (2017), the market consists of four participants, while in our case, it consists of eight participants. In the larger market, we can expect a lower rate of the equilibrium allocations reached, as it is enough for one participant in the group to deviate from truthful strategy in the consequential way to distort the whole allocation.

[44]Notice, however, that straightforward strategies may not be obviously dominant in the game induced by the pick-an-object mechanism $\mathbb{S}$. One can easily see, for example, that in the pick-an-object implementation of serial dictatorship used in our experimental section, the worst outcome that could come from following a straightforward strategy is typically worse than the best outcome that can come from a deviation from it in the first step of the mechanism.

	Direct	PAO	OSP	Direc=PAO p-value	Direct=OSP p-value	PAO=OSP p-value
TTC cyclic	5.9%	5.9%	n/a	1.00	n/a	n/a
TTC acyclic	22.6%	26.2%	12.6%	0.72	0.05	0.01
SD	38.1%	50%	83.%	0.16	0.00	0.00

Notes: All the p-values are for the two-sided Fisher exact test for equality of proportions of the equilibrium strategy allocations by treatments. test is performed on the allocation-level data

Table 5. Proportions of equilibrium allocations reached.

Table 5 presents the proportion of equilibrium allocations by treatments. Under TTC with cyclic priorities, only 5.9% of allocations in Direct and PAO were equilibrium allocations. This low rate is not surprising, given the rate of the truthful strategies manipulations, and the fact that even one consequential manipulation distorts the allocations. Under TTC with acyclic priorities, we observe that the OSP treatment has the lowest rate of equilibrium allocations. Despite the higher rate of truthful strategies, as every deviation from the truthful strategy more likely to be consequential for the allocation, it leads to only 12.6% of equilibrium allocations on average, which is significantly lower than in Direct and PAO treatments. Finally, under SD, OSP outperforms Direct and PAO mechanisms. Thus, we replicate the finding by Li (2017).

While the comparison of the rate of equilibrium allocations is useful for a smaller market, or under very low deviations from truthful strategies, we think the criterion is not very informative about the consequences of deviations in the case of a high rate of deviations from truthful strategies. Another approach would be to analyze the difference in the consequence of deviation from truthful strategies from an individual perspective by treatments. We define cost of deviation as the average difference between the payoffs of those who played truthfully and those who deviated from truthfulness, controlling for the role in the market.

Result 4 (Cost of deviation from truthful strategy):

(1) Under the TTC rule with cyclic priorities the average cost of deviation from the truthful strategy under Direct is 3.97 euros, under PAO 3.61 euros, with no significant difference between treatments.
(2) Under the TTC rule with acyclic priorities the average cost of deviation from the truthful strategy in Direct is 2.43 euros, in PAO 1.98 euros, and in OSP 3.36 euros with no significant difference between treatments.
(3) Under the SD rule the average cost of deviation from the truthful strategy in Direct is 2.53 euros, in PAO 1.55 euros, and in OSP 6.01 euros with all differences between treatments being statistically significant.

Support:

Table 6 presents the results of the OLS estimation of the effect of misreporting on the payoff of the subjects.

	(1) Payoff TTC cyclic	(2) Payoff TTC acyclic	(3) Payoff SD
Non-truthful strategy	-3.97***	-2.43***	-2.53***
	(.29)	(.29)	(.24)
Non-truthful in PAO	.38	.45	.98**
	(.39)	(.38)	(.36)
Non-truthful in OSP		-.93	-3.48***
		(.60)	(.76)
Dummies for each participant ID in each round	yes	yes	yes
Observations	1344	2800	2072
No. of clusters	24	37	37
R^2	.287	.370	.759
log(likelihood)	-4063.06	-7796.13	-5136.51

Notes: OLS regression. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are clustered at the level of matching groups and are presented in parentheses. Non-truthful is a dummy for not playing the truthful strategy. Non-truthful in PAO is the interaction of the Non-truthful dummy and the dummy for PAO treatment. Non-truthful in OSP is the interaction of the Non-truthful dummy and the dummy for OSP treatment.

TABLE 6. *OLS regression of payoff on the dummy for non-truthful strategy*

Each regression includes 56 dummies for each combination of ID and round, to account for the "role-specific" fixed effects, as the roles (a combination of preferences and priorities/scores) vary the prospects of earning high payoffs. Thus, the coefficient for the non-truthful dummy presents the average differences between subjects who play truthfully relative to subjects who play non-truthfully in the Direct mechanism, controlling for the role of the subjects. Under TTC with cyclic priorities (Model (1) of Table 6), the deviation from the truthful strategy, on average, leads to a loss of 3.97 euros in Direct (note that the maximum payoff for the allocation is 22 euros), while the deviations are 38 cents less costly in PAO, though the difference is not significant. Under TTC with acyclic priorities (Model (2) of Table 6), the deviation from the truthful strategy, on average, leads to a loss of 2.43 euros in Direct, while the deviations are 45 cents less costly in PAO, and 93 cents more costly in OSP, though the differences are not significant. Finally, under SD, (Model (3) of Table 6), the deviation from the truthful strategy, on average, leads to a loss of 2.53 euros in Direct, while the deviations are 98 cents less costly in PAO, and 3.48 euros more costly in OSP, with all differences being significant. The difference can be explained by the fact that skipping in PAO is less consequential, due to intermediate updates on which objects are left for allocation, which is not the case in Direct. The highest cost of deviations in OSP can be explained by the fact that all deviations from the truthful strategy are payoff-relevant in SD, as it is a unique equilibrium strategy in OSP, unlike in the other treatments.