Obvious Dominance and Random Priority

Marek Pycia and Peter Troyan*

January 2017

Abstract

In environments without transfers, such as refugee resettlement, school choice, organ transplants, course allocation, and voting, we show that Random Priority is the unique mechanism that is obviously strategy-proof, Pareto efficient, and symmetric; hence providing an explanation for the popularity of this mechanism. We also construct the full class of obviously strategy-proof mechanisms, and explain why some of them are more popular than others via a natural strengthening of obvious strategy-proofness.

1 Introduction

The central concerns in designing mechanisms for refugee resettlement, school choice, organ transplantation, course allocation, voting, and other social choice problems without transfers are participants' incentives as well as normative goals such as efficiency and fairness.[1] In particular, assuring that agents play the games correctly is crucial for attaining the normative properties of the mechanism in actual play. Dominant-strategy incentive compatibility, also known as strategy-proofness, assures that truth-telling is a dominant strategy in direct mechanisms, but this is useful only to the extent the participants understand it. Li (2016) proposes a refinement of strategy-proofness called *obvious strategy-proofness (OSP)*

*UCLA and University of Virginia, respectively. First presentation: April 2016. First posted draft: June 2016. For their comments, we would like to thank Itai Ashlagi, Yannai Gonczarowski, Ed Green, Fuhito Kojima, Shengwu Li, Giorgio Martini, Stephen Morris, Erling Skancke, Utku Unver, the Eco 514 students at Princeton, and the audiences at the NBER Market Design workshop and NEGT'16. Simon Lazarus provided excellent research assistance. Marek would like to acknowledge the financial support of the William S. Dietrich II Economic Theory Center at Princeton.

[1]See e.g. Roth (2015) and Jones and Teytelboym (2016) for resettlement, Abdulkadiroğlu and Sönmez (2003) for school choice, Roth, Sönmez, and Ünver (2004) for transplants, Sönmez and Ünver (2010) and Budish and Cantillon (2010) for course allocation, and Arrow (1963) for voting and social choice.

that characterizes the extensive-form mechanisms that are easy to understand as strategy-proof. While strong incentive properties are desirable, they can lead to tradeoffs on other dimensions, which raises the question which social choice functions are OSP implementable.

We show that there is an essentially unique mechanism that is obviously strategy-proof, efficient, and symmetric, and it is the well-known extensive-form Random Priority mechanism.[2] Thus, our results give an explanation for the popularity of this mechanism. We also construct the full class of obviously strategy-proof mechanisms for social choice environments without transfers; we call these mechanisms *millipede* games. While some millipede games such as sequential dictatorships are frequently encountered and they are indeed simple to play, others are rarely observed in market-design practice, and their strategy-proofness is not necessarily immediately clear. To further delineate the class of mechanism that are simple to play, we introduce a refinement of Li's concept, which we call *strong obvious strategy-proofness (SOSP)*. We show that strongly obviously strategy-proof and efficient mechanisms are almost sequential dictatorships.

An imperfect information extensive-form game is obviously strategy-proof if, whenever an agent is called to play, there is a strategy such that even the worst possible final outcome from following this strategy is at least as good as the best possible outcome from taking any other action at the node in question, where what is possible may depend on the actions of other agents in the future (and, in case of the deviation, it may also depend on how the agent plays following the deviation). Consider, for instance, object allocation via Random Priority, implemented as an extensive-form game in which nature draws the ordering in which agents move, and when given the opportunity to move, each agent can choose among the objects that remain after the earlier movers' choices. After nature's choice of ordering, this game takes the extensive form of a deterministic serial dictatorship, in which players take turns choosing their most preferred outcome. The game is hence obviously strategy-proof because each agent moves only once and can chose their outcome from a pre-determined set of available outcomes. While this extensive-form game is obviously strategy-proof, the corresponding direct mechanism is not.[3]

[2]By effectively unique, we mean that every OSP, efficient, and symmetric mechanism is equivalent to random priority.

[3]The observation that the extensive-form of random priority is obviously strategy-proof is due to Li (2016). Note also that, intuitively, describing a direct-revelation game in terms of an obviously strategy-proof game played by proxies makes the direct-revelation game simpler to play than alternative descriptions. Another extensive form game also called random priority is as follows: nature draws an agent uniformly at random, this agent picks an object, then nature draws an agent who has not moved yet, this agent picks an object, etc. These two games, while not identical, are equivalent in the following sense: for each profile of agents' preferences, each game admits a profile of obviously dominant strategies such that the two games lead to the same outcomes when the obviously dominant strategies are played. For simplicity, in this introduction we ignore differences between equivalent games.

Our first main result, Theorem 1, provides an explanation for the popularity of the Random Priority mechanism in allocation problems, both formal and informal, including house allocation, school choice, and course allocation, among others. Random Priority is obviously strategy-proof, as recognized by Li; it is also well-known to be efficient and to treat agents symmetrically.[4] Theorem 1 shows that no other mechanism satisfies these three properties: an extensive-form game is obviously strategy-proof, efficient, and symmetric if and only if it is random priority.[5] This insight resolves positively the quest to establish random priority as the unique mechanism with good incentive, efficiency, and fairness properties.[6]

The proof of Theorem 1 builds on our second main result, Theorem 2, which constructs the full class of OSP mechanisms as the class of *millipede* games. In a millipede game, first nature chooses a deterministic path, and then agents engage in a game of passing and clinching that resembles the well-studied centipede game (Rosenthal, 1981). To describe this deterministic subgame, we focus for simplicity on allocation problems with agents who demand at most one unit; the general case is similar. An agent is presented with some subset of objects that she can 'clinch', or, take immediately and leave the game; she also may be given the opportunity to 'pass', and remain in the game.[7] If this agent passes, another agent is presented with an analogous choice. Agents keep passing among themselves until one of them clinches some object and leaves the game (i.e., never moves again). When an agent clinches an object, this is her last move and she is going to be allocated this object at the end of the game.[8] Common examples of millipede games include Random Priority and Serial Dictatorships, which are millipede games in which the agent who moves can always clinch any object that is still available. However, these are not the only mechanisms that can be implemented as millipede games: the key restriction implied by obvious strategyproofness (and thus the defining feature of millipede games) is, once an agent is offered some object, if she passes, she is promised that later in the game, she will be able to clinch at least everything

[4]For discussion of efficiency and symmetry see, e.g., Abdulkadiroğlu and Sönmez (1998), Bogomolnaia and Moulin (2001), and Che and Kojima (2010).

[5]Efficiency assumes that agents play their dominant strategies, i.e., report truthfully. As noted above, if agents fail to do so efficiency may be lost; hence the importance of making allocation games simple to play.

[6]In single-unit demand allocation with at most three agents and three objects, Bogomolnaia and Moulin (2001) proved that random priority is the unique mechanism that is strategy-proof, efficient, and symmetric. In markets in which each object is represented by many copies, Liu and Pycia (2013) and Pycia (2011) proved that random priority is the asymptotically unique mechanism that is symmetric, asymptotically strategy-proof, asymptotically ordinally efficient (a more demanding efficiency criterion than ex post Pareto efficiency). While these earlier results looked at either very small or very large markets, our characterization holds for any number of agents and objects.

[7]In general, we may allow this agent to clinch the same object in several ways; while the choice between them has no impact on the agent's outcome, it might affect the allocation of others.

[8]In a centipede game, two agents pass back and forth until one of them 'takes' the pot and leaves the game. Our games have a similar take-or-pass structure, but with many more options of what to take; i.e., they look like centipede games with more "legs" (see Figure 4). Hence, we call these games *millipede* games.

she was able to clinch in the past, and possibly more. It is obvious that serial dictatorships satisfy this property; however, more complex mechanisms such as Gale's top trading cycles will not.[9] Other known mechanisms that can be implemented as millipede games (and hence, will be OSP) include bi-dictatorships, see e.g. Ehlers (2002). Our construction of the class of millipede games provides the precise characterization of the set of OSP implementable mechanisms.

The key assumption behind these results is that outcomes an agent is not indifferent between can be ranked in any order. This key assumption is trivially satisfied in standard models without transfers, but it excludes environments with transfers.[10] The results are robust to how agents' evaluate lotteries over outcomes, and they are derived in a very general model encompassing the standard allocation and social choice environments.

While millipede games are obviously strategy-proof, their play may still require substantial foresight on the part of the agents, similar to the foresight required in centipede games. For instance, at a node in a millipede game, a player might be offered the possibility of clinching his second-choice object, but not his top choice object even though it is still available. The obviously dominant strategy requires this agent to pass, but, if the agent passes, he might not be given the opportunity to clinch any of his top seventeen objects in the next one hundred moves. While this problem does not arise in the environment with transfers focused on in Li, it arises in the settings without transfers we study, and poses the question how to identify the extensive-form games that do not require such high degrees of future planning in environments without transfers.

We answer this question by introducing a refinement of obvious strategy-proofness, which we call *strong obvious strategy-proofness.* An extensive-form game is strongly obviously strategyproof if, whenever an agent is called to play, there is an action such that even the worst possible final outcome from that action is at least as good as the best possible outcome from any other action, where what is possible may depend on all future actions, including actions by the agent's future-self. We show that strong obvious strategy-proofness eliminates the complex members of the more general class of millipede games, and strongly obviously strategy-proof games take the simple form of almost sequential dictatorships (studied earlier by Pycia and Ünver (2016), in a different context).

Our paper builds on the key contributions of Li (2016) who formalized OSP and established its desirability as an incentive property by showing that in environments with transfers this property differentiates between ascending auctions and sealed-bid second-price

[9] See Li (2016) and Troyan (2016).

[10] Li (2016) constructed the class of obviously strategy-proof mechanisms in bilateral choice environments with transfers and quasilinear utilities, and he showed that in these environments the obviously strategy-proof mechanisms resemble ascending auctions.

auctions: though both are strategy-proof and revenue equivalent, ascending auctions are more popular than sealed bid auctions, and experimental subjects play dominant strategies in ascending formats significantly more often than in sealed-bid formats. Though his main focus is environments with transfers, without transfers, Li also shows that the extensive-form of random priority is obviously strategy-proof, while other standard mechanisms—top trading cycles and the static version random priority—are not.[11] We characterize the entire class of OSP mechanisms, and provide an explanation for the popularity of random priority over all other mechanisms, results which have no direct counterpart in his work. Finally, our analysis of strong obvious strategyproofness furthers our understanding of why some extensive forms of a mechanism are more often encountered in practice, despite both being obvious strategy-proof and equivalent from the persepctive of the Myerson-Riley revelation principle.[12]

Following up on Li's work, but preceding ours, Ashlagi and Gonczarowski (2016) study whether stable mechanisms such as Deferred Acceptance are obviously strategy-proof. They show that the answer is generally no, except in special environments where Deferred Acceptance simplifies to an obviously strategy-proof game with a 'clinch or pass' structure similar to millipede games. Troyan (2016) studies obviously strategyproof allocation via the popular Top Trading Cycles (TTC) mechanism, and provides a characterization of the priority structures under which TTC is OSP-implementable.[13]

More generally, this paper adds to our understanding of incentives, efficiency, and fairness in settings without transfers. In addition to Gibbard (1973, 1977) and Satterthwaite (1975), and the allocation papers mentioned above, the literature on mechanisms satisfying these key objectives includes Pápai (2000), Ehlers (2002) and Pycia and Ünver (2009) who characterized efficient and group strategy-proof mechanisms in settings with single-unit demand, and Pápai (2001) and Hatfield (2009) who provided such characterizations for settings with multi-unit demand.[14] Liu and Pycia (2013), Pycia (2011), Morrill (2014), and Hakimov and Kesten (2014) characterized mechanisms that satisfy incentive, efficiency, and fairness

[11] Li considers the original construction of top trading cycles of Shapley and Scarf (1974) in which each agent starts by owning exactly one object, and shows it is not obviously strategyproof. Of note is also Loertscher and Marx (2015) who study environments with transfers and construct a prior-free obviously strategy-proof mechanism that becomes asymptotically optimal as the number of buyers and sellers grows.

[12] For instance, our characterization of millipede games shows that there can be OSP mechanisms which are rarely seen in practice because they may require complex future planning on the part of the agents. While we think strong OSP is well-suited to differentiate between mechanisms in environments without transfers, it is is too strong in environments with transfers. For instance, fixed-price mechanisms satisfy SOSP, but ascending auctions in general do not.

[13] Following on our work, Bade and Gonczarowski (2016) study obviously strategy-proof and efficient social choice rules in several environments, and Arribillaga et al. (2016) study obviously strategy-proof voting rules.

[14] Pycia and Ünver (2016) characterized individually strategy-proof and Arrovian efficient mechanisms. For an analysis of these issues under additional feasibility constraints, see also Dur and Ünver (2015).

objectives.

2 Preliminaries

2.1 Model

Let $\mathcal{N}$ be a set of agents, and $\mathcal{X}$ a finite set of outcomes.[15] Each agent has a preference ranking over outcomes. The domain of preferences of agent $i \in \mathcal{N}$ is denoted $\mathcal{P}_i$, and a generic preference ranking in $\mathcal{P}_i$ is denoted by $\succ_i$. We allow for indifferences, and write $x \sim_i y$ if neither $x \succ_i y$ nor $y \succ_i x$. Every agent partitions the set of outcomes and is indifferent among all outcomes in the same element of the partition; in particular, each $\succ_i \in \mathcal{P}_i$ might be identified with a ranking of the elements of the partition. We denote by $I_i(x)$ the element of the partition of agent i that contains $x \in \mathcal{X}$. We assume that every strict ranking of the elements of the partition is in $\mathcal{P}_i$.[16]

This model encapsulates a wide variety of preference structures, including the set of all strict preferences and the set of all preferences. In addition, it includes allocation problems in which each agent cares only about his or her allocation. In the special case of object allocation, each agent only cares about her own assignment and is indifferent between how the remaining objects are assigned to others, and so each element of the partition of $\mathcal{X}$ can be identified with the allocation of agent i.[17] Where our assumptions fail is in settings with transfers in which each agent always prefers having more money to less.

When dealing with lotteries, we are agnostic as to how agents evaluate them, as long as the following property holds: an agent prefers lottery μ over ν if for any outcomes $x \in \operatorname{supp}(\mu)$ and $y \in \operatorname{supp}(\nu)$ this agent weakly prefers x over y; the preference between μ and ν is strict if, additionally, at least one of the preferences between $x \in \operatorname{supp}(\mu)$ and $y \in \operatorname{supp}(\nu)$ is strict. This mild assumption is satisfied for expected utility agents; it is also satisfied for agents who prefer μ to ν as soon as μ first-order stochastically dominates ν.

[15] The assumption that $\mathcal{X}$ is finite simplifies the exposition and it is satisfied in such examples of our setting as voting and the no-transfer allocation environments listed in the introduction. This assumption can be relaxed. For instance, our analysis goes through with no substantive changes if we allow infinite $\mathcal{X}$ endowed with a topology such that agents' preferences are continuous in this topology and the relevant sets of outcomes are compact.

[16] We slightly relax this assumption in Appendix D.

[17] In other words, in allocation problems $I_i(x)$ consists of all outcomes in which i obtains the same object as in x.

2.2 Obvious strategy-proofness

The standard notion of strategy-proofness is defined for normal form games, while obvious strategy-proofness applies to extensive-form games. An imperfect-information *extensive-form game* with perfect recall is defined in the standard way, as a collection of partially ordered *histories* (sequences of moves), where at every history h, one agent $i \in \mathcal{N}$ has a (finite) set of *actions* $A(h)$ from which to choose. We use the notation $h' \subseteq h$ to denote that h' is a subhistory of h (equivalently, h is a continuation history of h'), and write $h \subset h'$ when $h \subseteq h'$ but $h \neq h'$. An information set I of agent i is a set of histories such that for any $h, h' \in I$ and any subhistories $\tilde{h} \subseteq h$ and $\tilde{h}' \subseteq h'$ at which i moves at least one of the following two symmetric conditions obtains: either (i) there is a history $\tilde{h}^* \subseteq \tilde{h}$ such that $\tilde{h}^*$ and $\tilde{h}'$ are in the same information set, $A(\tilde{h}^*) = A(\tilde{h}')$, and i makes the same move at $\tilde{h}^*$ and $\tilde{h}'$, or (ii) there is a history $\tilde{h}^* \subseteq \tilde{h}'$ such that $\tilde{h}^*$ and $\tilde{h}$ are in the same information set, $A(\tilde{h}^*) = A(\tilde{h})$, and i makes the same move at $\tilde{h}^*$ and $\tilde{h}$. We denote by $I(h)$ the information set of the agent who moves at h.[18] Slightly abusing notation, we write (h, a) to denote the history h followed by the action $a \in A(h)$. Each terminal history is associated with an outcome in $\mathcal{X}$, and agents receive payoffs at each terminal history that are consistent with their preferences over outcomes $\succ_i$. While we do not consider incomplete information, our insights remain valid for situations in which agents' information is incomplete: we allow imperfect information, and the standard interpretation of incomplete information games as imperfect information games can be used.

A *strategy* for a player i is a function that specifies an action for agent i at each one of her information sets. A strategy profile $S = (S_i)_{i \in \mathcal{N}}$ is a list of strategies, one for each agent. When we want to refer to the strategy of a specific type $\succ_i$ of agent i, we will write $S_i(\succ_i)$. In addition, we allow for chance moves by nature. Let ω be a function specifying an action for nature at each history where nature is to move.[19] We sometimes write (S, ω) to denote a strategy profile together with nature's moves. An information set I is on path of the strategy profile S if there is a realization of nature's moves ω and a history $h \in I$ that is in the support of histories that can obtain given S. Given strategies S_{-i} of agents other than i, the first information set I at which two strategies S_i and S_i' diverge is an information set on path for agent i for both strategies and such that $S_i(I) \neq S_i'(I)$.[20]

[18]It will turn out to be without loss of generality to assume all information sets are singletons, and so we will be able to drop the $I(h)$ notation and identify each information set with the unique history taken to reach it.

[19]ω can be thought of as describing nature's "strategy", i.e., it is one particular realization of chance moves throughout the game at every history at which nature is called to play.

[20]Li (2016) refers to such an information set as an *earliest point of departure*. Note that for two strategies, there can be multiple earliest points of departure.

Following Li (2016), for a game Γ, a strategy S_i *obviously dominates* another strategy S_i' for player i if, at any first information set at which these two strategies diverge, the worst possible payoff to the agent from playing S_i is at least as good as the best possible payoff from the other strategy, ranging over all possible (S_{-i}, ω). A profile $(S_i(\cdot))_{i\in\mathcal{N}}$ of strategies is *obviously dominant* if for any player i and type $\succ_i$ the strategy $S_i(\succ_i)$ obviously dominates any other strategy. When there exists a profile of strategies $(S_i(\cdot))_{i\in\mathcal{N}}$ that is obviously dominant, we say Γ is *obviously strategy-proof*. An example of an obviously strategy-proof game is the well-known random priority game (see the definition and discussion in the introduction).[21]

An *extensive-form mechanism*, or simply a *mechanism*, is an extensive-form game Γ together with a profile of strategies $(S_i(\succ_i))_{i\in\mathcal{N}}$ for each preference (or type) profile $\succ$. Following Li, we assume that Γ is pruned, that is each action is on the path of the play of the obviously dominant strategies $(S_i(\succ_i))_{i\in\mathcal{N}}$ for some type profile. A mechanism is *obviously strategy-proof* if there exists a profile of strategies $(S_i(\succ_i))_{i\in\mathcal{N}}$ that are obviously dominant. A mechanism is *Pareto efficient* if, for all type profiles $(\succ_i)_{i\in\mathcal{N}}$, when the agents follow strategies $(S_i(\succ_i))_{i\in\mathcal{N}}$, every possible outcome (for all possible chance moves) is Pareto efficient, i.e., for all possible outcomes x, there is no other y such that $y \succeq_i x$ for all i, and $y \succ_i x$ for some i.

Two obviously strategy-proof extensive-form mechanisms $\left(\Gamma, (S_i(\cdot))_{i\in\mathcal{N}}\right)$ and $\left(\Gamma', (S_i'(\cdot))_{i\in\mathcal{N}}\right)$ are *equivalent* if, for every profile of types $(\succ_i)_{i\in\mathcal{N}}$, the distribution of outcomes in Γ when agents play $(S_i(\cdot))_{i\in\mathcal{N}}$ is the same as in Γ' when agents play $(S_i'(\cdot))_{i\in\mathcal{N}}$. Two games Γ and Γ' are *equivalent* if there are profiles of strategies $(S_i(\cdot))_{i\in\mathcal{N}}$ and $(S_i'(\cdot))_{i\in\mathcal{N}}$ such that mechanisms $\left(\Gamma, (S_i(\cdot))_{i\in\mathcal{N}}\right)$ and $\left(\Gamma', (S_i'(\cdot))_{i\in\mathcal{N}}\right)$ are obviously strategy-proof and equivalent. Lemma 2 shows that for each obviously strategy-proof game, there is an equivalent obviously strategy-proof game that has perfect information.

3 Random Priority

The random priority mechanism we are interested in is widely used in allocation settings, defined as follows. There is a set of objects $\mathcal{O}$, each object $o \in \mathcal{O}$ is represented by $|o|$ copies, where every $|o|$ is a positive integer.[22] The set of feasible allocations of an one agent is $Q \subseteq \times_{o\in\mathcal{O}} \{0, 1, ..., |o|\}$. We assume that if $q_i \in Q$ then Q contains all $\left(q^1, ..., q^{|\mathcal{O}|}\right) \neq (0, ..., 0)$ such that for each coordinate j we have $q^j \leq q_i^j$. For example, (i) in the school choice problem each agent demands at most one object and $Q = \left\{q \in \times_{o\in\mathcal{O}} \{0, 1, ..., |o|\} \mid \sum_{o\in\mathcal{O}} q^o \leq 1\right\}$; (ii)

[21] We consider pure strategies, but the analysis can be extended to mixed strategies.

[22] The set of objects may include a special object called the outside option, but its presence play no role in the main-text results. Appendix D discusses how, by explicitly taking outside options into account, we can slightly extend our results beyond the class of preference domains studied in the main text.

in house allocation problems without outside options, each agent demands exactly one object, $Q = \left\{ q \in \times_{o\in\mathcal{O}} \{0, 1, ..., |o|\} \mid \sum_{o\in\mathcal{O}} q^o = 1 \right\}$. Each outcome $x \in \mathcal{X}$ consists of a profile of allocations $q_i \in Q$ for agents $i \in \mathcal{N}$ such that $\sum_{i\in\mathcal{N}} q_{i,o} \leq |o|$ for each $o \in \mathcal{O}$. We assume that $\mathcal{X}$ is nonempty. Each agent has strict preferences over allocations from Q_i and each agent's preferences over outcomes are determined by this agent's allocation; that is each agent is indifferent among allocations of other agents. We assume that all strict preference rankings over an agent's allocations are in this agent's domain of preference rankings; hence this allocation model is a special case of our general model.[23]

Our main result in this setting is showing that the popular extended-form random priority mechanism—described in the Introduction—is characterized by obvious strategy-proofness together with Pareto efficiency and symmetry. A game Γ is symmetric if for all agents i and j game Γ is isomorphic to a game Γ' with the roles of i and j exchanged. Exchanging the roles of i and j in game Γ creates game Γ' such that there is a bijection σ of histories between Γ and Γ' satisfying the following properties: terminal histories are mapped into terminal histories; if h is a subhistory of H in Γ then $\sigma(h)$ is a subhistory of $\sigma(H)$ in Γ'; and whenever i moved at history h in Γ then j moves at history $\sigma(h)$ in Γ', and vice versa; whenever an agent $k \neq i, j$ moved at h in Γ then the same agent moves at $\sigma(h)$ in Γ'; and the outcome of i after terminal history h in Γ is the same as the outcome of j after terminal history $\sigma(h)$ in Γ', and vice versa, while payoffs of other agents at terminal history h is the same as at $\sigma(h)$.[24]

A mechanism is symmetric if the underlying game is symmetric. Symmetry and efficiency are well-known to be properties of random priority, and Li (2016) shows that random priority is obviously strategy-proof.

Theorem 1. *An extensive-form mechanism is obviously strategy-proof, symmetric, and efficient if and only if it is equivalent to random priority.*

This result explains why random priority is a very popular allocation mechanism: not only it is efficient, fair, and simple to play thanks to its obvious strategy-proofness; it is the only mechanism with these properties.

We prove this theorem in the appendix. The proof relies on our characterization of obviously strategy-proof mechanisms, to which we turn next. Knowing the structure of these mechanisms allows us to extend the bijective approach of Abdulkadiroğlu and Sönmez (1998) (cf. also Pathak and Sethuraman (2010) and Carroll (2010)) to our setting.

[23] In Appendix D, we relax the assumption that all strict preference rankings are in the preference domain.

[24] More generally, take a permutation of agents σ, and game Γ, and and define game $\sigma(\Gamma)$ recursively as follows: at every node at which nature moves, it has same moves in $\sigma(\Gamma)$ as in Γ; at every history at which i moves in Γ agent $\sigma(i)$ moves in $\sigma(\Gamma)$ and has the same moves as i in Γ, at every terminal history the outcome of i in Γ is the same as the outcome of $\sigma(i)$ in the corresponding terminal history of $\sigma(\Gamma')$.

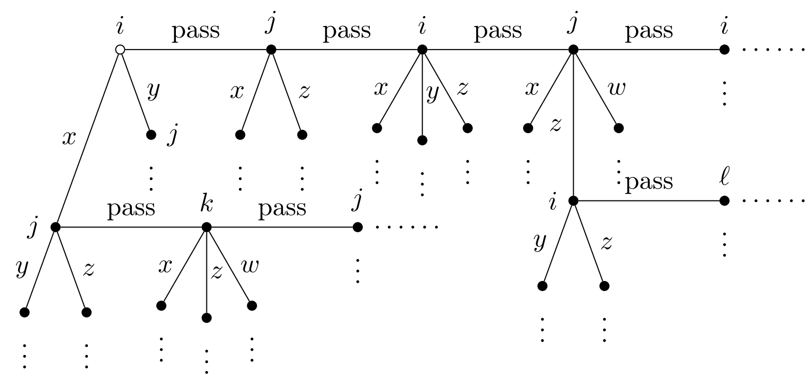

Figure 1: An example of a millipede game in the context of object allocation.

4 Obvious Dominance and Millipede Games

In this section, we characterize the class of obviously strategy-proof mechanisms as a class of games that we call *millipede games*. Intuitively, a millipede game is a take-or-pass game similar to a centipede game, but with more players and more actions (i.e., "legs") at each node. Figure 4 shows the extensive form of a millipede game in an object allocation environment, where the agents are labeled $i, j, k, \ldots$ and the objects are labeled $w, x, y, \ldots$. At the start of the game, the first mover, agent i has three options: he can take x, take y, or pass to agent j.[25] If he takes an object, he effectively leaves the game, and the game continues with a new agent. If he passes, then agent j can take x, take z, or pass back to i. If he passes back to i, then i's possible choices increase from his previous move (he can now take z). The game continues in this manner until all objects have been allocated.

While Figure 4 considers an object allocation environment, millipede games can be defined more generally on any preference structure that satisfies our assumptions of Section 2. Recall that we use the term *payoff* to refer to the indifference class in the partition on an agent's preference domain. In a complete-information extensive-form game, we say that a payoff x is *possible* for agent i at history h if there is a strategy profile of all the agents such that h is on the path of the game and agent i obtains payoff x; we denote by $P_i(h)$ the set of payoffs that are possible for i at h. At a history h, if, by taking some action $a \in A(h)$ an agent never moves again in the game and receives the same payoff for every terminal $\bar{h} \supseteq (h, a)$, we say that a is a *clinching* action; we denote by $C_i(h)$ the set of payoffs i can clinch at h.[26] If an action is not a clinching action, then it is called a *passing* action. For

[25]In the general definition of a millipede game, it will be possible that none of the actions are passing actions and so all actions are taking actions. However, if there is a passing action, there can only be one, to ensure obvious strategyproofness. This will be explained in detail below.

[26]Clinching actions are generalizations of the "taking actions" of Figure 4 to environments where the outcomes may be different from object allocation.

any history h of a game we denote by $\mathcal{H}_i(h) = \{h' \subsetneq h : i \text{ is the mover at } h'\}$ the set of all predecessor histories at which i moves (not including h itself).

A **millipede game** is a finite imperfect information extensive form game that satisfies the following properties. Nature either moves once, at the empty history $\emptyset$, or the game has perfect information and nature has no moves.[27] At any history h at which an agent, say i, moves, all but at most one action are clinching actions; the remaining action (if there is one) is a passing action. Each clinching action is associated with a unique payoff of agent i; if i chooses a clinching action, he does not move again, and he ultimately receives the payoff associated with the clinching action chosen.[28] Furthermore, for all payoffs $z \in \mathcal{X}$, at least one of the following holds: either (a) $z \in P_i(h)$; or (b) $z \notin P_i(\tilde{h})$ for some $\tilde{h} \in \mathcal{H}_i(h)$; or (c) $z \in \cup_{\tilde{h} \in \mathcal{H}_i(h)} C_i(\tilde{h})$; or (d) $\cup_{\tilde{h} \in \mathcal{H}_i(h)} C_i(\tilde{h}) \subseteq C_i(h)$.

Notice that that millipedes have recursive structure: the continuation game that follows any action is also a millipede game. A simple example of a millipede game is a sequential dictatorship in which no agent ever passes and there is only one active agent at each node.[29] A more complex example is given in Figure 4.[30] We provide an alternative definition of millipede games in the appendix.

Theorem 2. *A game Γ is obviously strategy-proof if and only if it is equivalent to a millipede game.*

This theorem is applicable in many environments, including allocation environments in which agents care only about the object(s) they receive; in this case a payoff is simply the object(s) an agent receives. In this paper we focus attention on these allocation environments. Another environment Theorem 2 applies to is the standard social choice environment in which no agent is indifferent between any two outcomes. Theorem 2 has straightforward corollaries for deterministic mechanisms: in deterministic environments, we just skip nature's move, and continue with a standard millipede game.

Notice that in voting environments in which all agents can strictly rank all outcomes, Theorem 2 implies that each OSP game is equivalent to a game in which either there are only two outcomes that are possible when the first agent moves (and in such case, the game is equivalent to a game in which the first agent to move either can clinch any of them, or

[27] We allow games in which an agent has only one action.

[28] There may be several clinching moves associated with the same payoff.

[29] We call an agent *active* at history h of a millipede game if the agent moves at h, or the agent moved prior to history h and has not yet clinched an outcome.

[30] The first more complex example of a millipede game we know of is due to Ashlagi and Gonczarowski (2016). They construct an example of OSP-implementation of deferred acceptance on some restricted preference domains (though they do not classify the actions as "passing" or "clinching" actions). On these restricted domains, DA reduces to a millipede game.

can clinch one of them or pass to the other agent who can then clinch either of the two outcomes), or the first agent to move can clinch any possible outcome and has no passing action. The latter case is the standard dictatorship, with a possible restricted set of possible outcomes, while the former case resembles the almost-sequential dictatorships we study in the next section.[31]

One of the implications of Theorem 2 follows through straightforward recursion: every millipede game is obviously strategy-proof when the agents play the following strategies, which we refer to as *greedy strategies:* at each move at which the agent can clinch the best still possible outcome for her, the strategy has the agent clinch this outcome; otherwise, the agent passes. We provide the details and the rest of the proof of Theorem 2 in the appendix, and explain the key steps of the other direction below.

Step 1. *Every OSP game Γ is equivalent to a perfect information OSP game Γ' in which nature moves once, as the first mover.*

This follows because if we break any information set with imperfect information to several different information sets with perfect information, the set of outcomes that are possible shrinks. For an action a to be obviously dominant, the worst possible outcome from a must be (weakly) better than the best possible outcome from any other a'. If the set of possibilities shrinks (in the set inclusion sense), then the worst case from a only improves, and the best case from a' worsens; thus, if a was obviously dominant in Γ, it will remain so in Γ'.[32]

Step 2. *At every history, all actions except for possibly one are clinching actions.*

Step 2 allows us to greatly simplify the class of OSP games to "clinch or pass" games. An action *clinches* x if for every possible history, agent i receives some y such that $y \sim_i x$.[33] In addition to the clinching actions, there may be at most one action that does not clinch an outcome for i, which we call a *passing* action. The main insight here is that there cannot be more than one such action. Indeed, if there were two such actions a and a', then following each of a and a' there are at least two outcomes that are possible, but not guaranteeable.

[31] If the game is efficient in addition to being OSP, then the sets of outcomes the agents' choose from are not constrained except possibly for the case when there are two outcomes in the economy, and the first agent can clinch one of them or pass the decision to the other agent (who can then clinch either outcome).

[32] That every OSP game is equivalent to an OSP game with perfect information was first pointed out in a footnote by Ashlagi and Gonczarowski (2016). They also make the point that de-randomizing an OSP game leads to an OSP game.

[33] The complete argument is more subtle, because it may be that following action a, several different payoffs are possible, but agent i is able to guarantee each one (using different strategies in the future). We say that a payoff x is *guaranteeable* for agent i at history h if there is a strategy of i such that (i) there is a strategy of other agents such that h is on the path of the game, and (ii) for any profile of the strategies of other agents such that h is on the path of the game, agent i obtains payoff x. One of the main lemmas proved in the appendix simplifies the analysis by showing that for any OSP game Γ, there is an equivalent OSP game Γ' such that all actions except possibly one are clinching actions.

Thus, it will always be possible to find a type of agent i for which one of the possibilities following a is at best his second choice, while one of the possibilities following a' is his first choice, which implies that a does not obviously dominate a'.

While step 2 shows that for almost all actions i's outcome is completely determined, what is perhaps more interesting is that there even *can* be one action such that i's outcome is uncertain without violating OSP. The reason for this follows from step 3.

Step 3. *If agent i passes at a history h, then the payoff she ultimately receives must be at least as good as any of the payoffs she could have clinched at h.*

An agent may follow the passing action if she cannot clinch her favorite possible outcome today, and so she passes, hoping she will be able to move again in the future and get it then. To retain obvious strategy-proofness the game needs to promise agent i that she can never be made worse off by passing: she will at least be able to clinch every payoff that she could have clinched in the past, and may possibly be able to clinch even more. This is the reason one of the conditions (a)-(d) obtains.

Combining steps 1-3 imply that any OSP game Γ is equivalent to a millipede game.

5 Strong Obvious Strategy-Proofness

Even mechanisms that are OSP and efficient can still be quite complex to actually play. The reason is that, while OSP assumes that agents do not understand how the choices of *other* agents will translate into outcomes, it still presumes that they understand how their own future actions affect outcomes. Thus, while OSP guarantees that agents do not have to reason carefully about others, it still requires that they do so with regard to their own "future self". Thus, an obviously strategy-proof mechanism may not be simple enough for agents to play correctly. Here, we introduce a strengthening that we call *strong obvious strategy-proofness* (SOSP).

Definition 1. Strategy S_i *strongly obviously dominates* strategy S_i' for agent i with preferences θ_i if at the earliest point (information set) of departure I between $S_i'(\theta_i)$ and $S_i(\theta_i)$ the best possible outcome from playing $S_i'(\theta_i)$, across all strategies and types of other agents, is at most as good for agent i of type θ_i as the worst outcome from playing $S_i(\theta_i)$ at information set I followed by by any play by other agents and any play by agent i.[34] If a strategy S_i strongly obviously dominates all other S_i', then we say that S_i is *strongly obviously dominant.*

Similarly to Li's OSP, the above stronger concept corresponds to a class of agents cognitive limitations. For simplicity, let us consider perfect-information games. We say that

[34]For two strategies S_i and S_i', an information set I is called an *earliest point of departure* if $S_i(I') = S_i'(I')$ for all information sets I' such that there are $h' \in I$ and $h \in I$ such that $h' \subset h$, but $S_i(I) \neq S_i'(I)$.

an extensive-form game Γ is outcome-set equivalent to extensive-form game Γ' if there is a bijection ϕ between histories such that the set of possible outcomes after history h is the same as the set of possible outcomes after $\phi(h)$. We then obtain.[35]

Theorem 3. *A game Γ and strategy profile S are SOSP if and only if the corresponding strategy profile S' is strategy-proof in any outcome-set equivalent game Γ'.*

While in the environments with transfers studied by Li, the canonical obviously strategy-proof mechanisms—ascending auctions—fail to satisfy this stronger concept, in environment without transfers, the canonical obviously strategy-proof mechanisms—sequential dictatorships—do satisfy SOSP.

The SOSP requirement basically characterizes sequential dictatorships. We say that a mechanism is a *curated dictatorship* if it is a perfect-information game in which nature moves first, and then the agents move in turn, with each agent moving at most once. The ordering of the agents and the sets of payoffs from which they choose is determined by nature's move, and the moves of earlier agents. As long as there are at least three payoffs possible for an agent who moves, at his move, this agent can clinch any of the possible payoffs (while clinching any of the payoffs the agents also picks a message from a pre-determined set of messages). When only two payoffs are possible, the agent can be faced with either a choice between them (including picking an accompanying message), or, he might be given a possibility to clinch one of these objects (and picking an accompanying message) and passing (with no message).

Theorem 2 then implies the following.[36]

Corollary 1. *A mechanism is SOSP if and only if it is equivalent to a curated dictatorship.*

We further say that a mechanism is an *almost-sequential dictatorship* if it is a perfect-information game in which nature moves first; then the agents move in turn (the ordering being determined by nature's move, and the moves of earlier agents) with each agent moving at most once. At his move, an agent picks his objects and sends a message. As long as there are at least three objects unallocated, the moving agent can choose from all still available

[35]While Li also shows that OSP mechanisms are precisely the mechanisms that can be implemented with bilateral commitment, this result does not extend to our setting. The reason is that bilateral commitment presumes that agents are perfectly forward looking and do not make errors in single-agent games, and SOSP relaxes this assumption.

[36]To prove this corollary it is sufficient to notice that pruned millipedes in which an agent has three or more possible payoffs and a passing move are not SOSP. Suppose not and notice that we can assume that the agent has at least one clinching move (clinching payoff x) and a passing move. Furthermore, among outcomes possible after passing there is an outcome y that cannot be clinched at this move, and a third outcome $z \neq x, y$. Then, an agent with preference ranking in which $y \succ x \succ z \succ \ldots$ has no strongly obviously dominant action.

objects. When there are two objects remaining, an agent can be faced with either a choice between them, or, he might be given a choice between one of these object for sure or giving the next agent an opportunity to allocate the remaining two objects among the two of them.[37]

Theorem 4. *A game and strategy-profile are SOSP and efficient if and only if the game is equivalent to an almost-sequential dictatorship.*

We are in the process of planning an experiment checking whether SOSP extensive form games are indeed simpler to play that equivalent OSP but not SOSP games. We conjecture that the experiment will confirm the natural intuition that subjects will find it easier to play an almost-sequential dictatorship than to play a more complex millipede game.

6 Conclusion

Obvious strategy-proofness is a highly desirable incentive property, as convincingly argued by Li (2016). The present paper follows up on Li's work by analyzing obvious strategy-proofness in environments without transfers. The paper's main insight is that random priority is the unique obviously strategy-proof mechanism that satisfies basic efficiency and fairness properties, thus providing an explanation why random priority is popular in practical allocation problems. We also characterize obviously strategy-proof games.[38] Arguing that playing an obviously dominant strategy might require extensive foresight in environments without transfers, we propose and analyze the more demanding concept of strong obvious strategy-proofness; games that are strongly obviously strategy-proof are indeed simple to play.

A Proof of Theorem 2

The key to understanding our analysis of OSP games and the proof of Theorem 2 are the concepts of possible, guaranteeable, and clinchable outcomes that we now introduce more formally. In an extensive-form game, every time an agent takes an action, she shrinks the set of outcomes (terminal histories) that can obtain. Determining whether an action is obviously dominant at a given history requires comparing the set of possible payoffs following each action with the set of payoffs that can be guaranteed.

[37] Pycia and Unver (2016) use the same name for deterministic mechanisms without messages that belong to the class we study; they show that this are exactly the deterministic mechanisms which are strategy-proof and Arrovian efficient with respect to a complete social welfare function. We use the name their introduced because our class is a natural extension of theirs.

[38] While we focus on analysis obviously strategy-proof games, our analysis can be rephrased in terms of social choice rules.

Formally, for a game Γ, let $z(h, S, \omega)$ be the final outcome associated with the terminal history that is reached when starting from history h and play proceeding according to the strategy profile (S, ω). Then, let

$$X_i(h, S_i) = \{x \in \mathcal{X} : z(h, S_i, S_{-i}, \omega) = x \text{ for some } (S_{-i}, \omega)\}$$

denote the set of outcomes that are possible starting from history h if i follows strategy S_i, ranging over all possible strategies of the other agents and nature. Note that the agent acting at h may be some $j \neq i$, but we can still define the set of outcomes that are possible for i starting from this history if she follows strategy S_i at all future histories.

We first discuss the distinction between two types of payoffs (possible vs. guaranteeable) and then the distinction between two types of actions (clinching vs. passing). Consider an agent i of type $\succ_i$ and an indifference class $I_i(x)$. If there exists some S_i such that $y \in X_i(h, S_i)$ for some $y \sim_i x$, then we then we say that payoff x is **possible** for i at h. If, further, there exists some S_i such that $y \sim_i x$ for all $x, y \in X_i(h, S_i)$, then we say payoff x is **guaranteeable** for i at h. Note that we are slightly abusing notation, since when we say x is possible at h, what we mean is some outcome in the indifference class $I_i(x)$ is possible at h. Similarly, when we say x is guaranteeable for i at h, we mean that there is some strategy such that agent i can ensure that some outcome in the indifference class $I_i(x)$ will obtain, no matter the actions of the other agents or nature.[39] Let

$$P_i(h) = \{x \in \mathcal{X} : \exists S_i \text{ s.t. } y \sim_i x \text{ for some } y \in X_i(h, S_i)\}$$

$$G_i(h) = \{x \in \mathcal{X} : \exists S_i \text{ s.t. } y \sim_i x \text{ for all } y \in X_i(h, S_i)\}$$

be the sets of payoffs that are possible and guaranteeable at h, respectively. Note that $G_i(h) \subseteq P_i(h)$, and the set $P_i(h) \setminus G_i(h)$ is the set of payoffs that are possible at h, but are not guaranteeable at h. This latter set will be key to many of our arguments.

Last, we define a distinction between two kinds of actions: clinching actions and passing actions. Let i be the agent who is to act at a history h, and define $poss_i(a) = P_i((h, a))$ to be the payoffs that are possible for i if she follows action a at h. If for all $y \in P_i((h, a))$, we have $y \sim_i x$, then we say that action $a \in A(h)$ **clinches** payoff x for i. If, at history h, there exists an action a that clinches x for i, we say that x is clinchable at h, and call a a clinching

[39] For example, in object allocation, $x \in \mathcal{X}$ denotes the entire allocation of all agents. If i is the first mover in a serial dictatorship and she takes some object o, formally there are still many outcomes $x \in \mathcal{X}$ for which she receives o. Since she is indifferent between all outcomes where she receives o, we say that the payoff x is guaranteeable.

action. In words, following a clinching action, i's outcome is completely determined (modulo indifference classes). Note that there can be more than one action that clinches the same payoff x, and it is possible that i is called to move again following a clinching action a. Any action of an agent that is not clinching is called a **passing action**.

For clarity, the proof is organized according to the three steps described in the main text. Each step is formalized and proved using one or more lemmas.

Step 1: Every OSP game Γ is equivalent to a perfect information OSP game Γ' in which nature moves once, as the first mover.

We break this result into two lemmas. The first shows that we can assume perfect information, and the second shows that we can assume nature moves only once.

Lemma 1. *(Ashlagi and Gonczarowski, 2016) Every OSP game is equivalent to an OSP game with perfect information.*

Proof. Ashlagi and Gonczarowski (2016) mention this result in a footnote; here, we provide the straightforward proof for completeness. Denote by $A(I)$ the set of actions available at information set I to the agent who moves at I. Take an obviously strategy-proof game Γ and consider its perfect-information counterpart Γ', that is the perfect information game at which at every history h in Γ the moving agent's information set is $\{h\}$ in Γ', the available actions are $A(I)$, and the outcomes in Γ' following any terminal history are the same as in Γ. Notice that the support of possible outcomes at any history h in Γ' is a subset of the support of possible outcomes at $I(h)$ in Γ. Hence, Γ' is obviously strategy-proof and equivalent to Γ'. $\square$

Lemma 2. *Every OSP game is equivalent to a perfect information OSP game in which nature moves once, as the first mover.*

Proof. This lemma can be derived from the comment in Ashlagi and Gonczarowski (2016) that each de-randomized OSP mechanism is OSP. We provide the straightforward direct proof. In light of the transitivity of the equivalence relationship, the previous lemma tells us that we can assume that Γ has perfect information. We can further assume that nature moves at the empty history $\emptyset$ (possibly just having one move at this history). Let $\mathcal{H}_{\text{nature}}$ be the set of histories h at which nature moves. Consider a modified game Γ' in which at the empty history nature chooses actions from $\times_{h\in\mathcal{H}_{\text{nature}}} A(h)$ and at $h \in \mathcal{H}_{\text{nature}} - \{\emptyset\}$ nature chooses from a singleton set of actions that contains the choice from $A(h)$ that nature made at $\emptyset$. Then, the support of possible outcomes at any history h in Γ' is a subset of the support of possible outcomes at the corresponding history in Γ, where the corresponding histories are defined by mapping the $A(h)$ component of the action taken at $\emptyset$ by nature in Γ' as an action made by nature at h in game Γ. Hence, Γ' is obviously strategy-proof, and Γ and Γ' are equivalent. $\square$

Step 2: At every history, all actions except for possibly one are clinching actions.

Lemma 3. In any obviously strategyproof pruned game Γ, at any history h there is at most one action $a^* \in A(h)$ such that $P_i((h,a^*)) \not\subseteq G_i(h)$.

Proof. By way of contradiction, suppose at h there are two passing actions. In particular, there is an agent i who moves at h.

Step 1. Suppose there are two distinct payoffs $x, y \in P_i(h) \setminus G_i(h)$ and preference type $\succ_i$ such that (i) x and y are the first and second $\succ_i$ −best possible payoffs in $P_i(h)$, and (ii) h is on the path of the game for type $\succ_i$ when all agents follow the greedy strategy. Then, there is at most one action $a^* \in A(h)$ such that $P_i((h,a^*)) \not\subseteq G_i(h)$, and the greedy strategy of type $\succ_i$ chooses action a^* at h.

To prove the claim of this step, it is enough to consider the case $x \succ_i y$. If $x \in poss_i(a)$ for some $a \in A(h)$, then $y \in poss_i(a)$ as well. Indeed, if not then $x \in poss_i(a)$ and $y \notin poss_i(a)$. For type $\succ_i$, the worst case outcome from following a is strictly worse than y because x and y are assumed to be the $\succ_i$ −best possible payoffs in $P_i(h)$, x is not guaranteeable at h, and y is not possible following a. Because y is possible, action a is not obviously dominant. Furthermore, there is some $a' \in A(h)$ such that $a' \neq a$ because $y \in P_i(h) \setminus G_i(h)$. As x is not guaranteeable at h, the worst case outcome from any such a' is strictly worse than x, while the best case outcome from a is x, and so no a' can obviously dominate a_1. Thus, type $\succ_i$ has no obviously dominant action, which contradicts that the game is OSP, and proves the claim of this paragraph.

To prove the claim of Step 1 by way of contradiction, assume that there are two actions a_1^* and a_2^* that could lead to possible but not guaranteeable payoffs. Consider some $x \in P_i(h) \setminus G_i(h)$. By the previous paragraph, we know that $x \in poss_i(a_1^*)$ and $x \in poss(a_2^*)$. However, by assumption, x is not guaranteeable at h, and so, for type $\succ_i$, the worst case payoff from *any* action a' must be strictly worse than x, while the best case outcomes from a_1^* and a_2^* are both x. Therefore, no action $a' \in A(h)$ is obviously dominant, which contradicts that Γ is OSP.

We can conclude that there is at most one action a^* that leads to x for some continuation strategies of players. Because x is only possible following a^*, any obviously dominant strategy of any type $\succ_i$ that ranks any $x \in P_i(h) \setminus G_i(h)$ first among the payoffs in $P_i(h)$ must select a^* at this history, concluding the proof of the claim of Step 1.

Step 2. Suppose i moves at history h. Then, there is at most one action $a^* \in A(h)$ such that $P_i((h, a^*)) \not\subseteq G_i(h)$.

To prove the claim of this step, first consider any earliest history h_0^i at which i is to move, and choose some $x \in P_i(h_0^i) \setminus G_i(h_0^i)$. If there are two payoffs $x, y \in P_i(h_0^i) \setminus G_i(h_0^i)$, then we can apply Step 1 to type $x \succ_i y \succ_i \cdots$. If x is the only payoff in $P_i(h_0^i) \setminus G_i(h_0^i)$, the conclusion follows immediately. Indeed, assume there were two, a_1^* and a_2^*. Then, the worst case from any action is strictly worse than x (because x is not guaranteeable), while the best case from both a_1^* and a_2^* is x, so nothing can obviously dominate a_1^*, and a_1^* does not obviously dominate a_2^*.

Now, consider any successor history $h' \supset h_0^i$ at which i is to move, and make the inductive assumption that at every $h \subset h'$, there is only one possible passing action. First, consider the case where $|P_i(h') \setminus G_i(h')| = 1$, and let x be the unique payoff that is possible but not guaranteeable. By way of contradiction, assume there were two actions, a_1 and a_2, such that x was a possible outcome. Some type $\succ_i$ must be receiving x at some terminal history $\bar{h} \supset h'$. If x is the top choice among all payoffs in $P_i(h')$ for some type $\succ_i$ for which history h' is on-path, then neither a_1 nor a_2 can be obviously dominant.[40] If x is not the best possible payoff in $P_i(h')$ for any such type, then, since x is the only payoff that is possible, but not guaranteeable, every other payoff in $P_i(h')$ is guaranteeable. This implies that at h', each relevant type of agent i can guarantee herself best possible outcome in $P_i(h')$, and so no type should play strategies a_1 nor a_2; a contradiction because the game is pruned.

Finally, assume $|P_i(h') \setminus G_i(h')| \geq 2$, and let $x, y \in P_i(h') \setminus G_i(h')$. First, consider the case that there is some $x \in P_i(h') \setminus G_i(h')$ such that $x \notin G_i(h)$ for all $h \subset h'$ at which i is to move. Recall the inductive hypothesis says that at all such h, there is a unique passing action. Thus, if x has never been guaranteeable, all types $x \succ_i \cdots$ must have followed the unique passing action at all such $h \subset h'$, and we can apply Step 1 to the type $x \succ_i y \succ_i \cdots$. The last case to consider is where all $x, y \in P_i(h') \setminus G_i(h')$ were also guaranteeable at some earlier history h. Consider a type $\succ_i$ of agent i who receives payoff x at some terminal history $\bar{h} \supset h'$. First, note that for any $z \succ_i x$, $z \notin G_i(h')$ as otherwise this type would not follow a strategy whereby x was a possible outcome.[41] Let z be the $\succ_i -$best payoff in $P_i(h')$, and w be the second-best possible payoff in $P_i(h')$ for this type; we allow $w = x$. We can again apply Step 1 to type $\succ_i$ and conclude there is at most one passing action a^*. QED

Lemma 4. *Let Γ be an obviously strategyproof game that is pruned with respect to the strategy*

[40] Since x is not guaranteeable, the worst case outcome from a_1 is strictly worse than x, while the best case outcome from a_2 is x, so a_1 does not obviously dominate a_2. An equivalent argument shows a_2 does not obviously dominate a_1.

[41] If x is the $\succ_i$-best possible payoff in $P_i(h')$ for all types that reach h', then apply the same argument to a type that receives y at some terminal history and set $z = x$.

profile $(S_i(\succ_i))_{i\in\mathcal{N}}$*. There exists an equivalent obviously strategyproof game* Γ' *with perfect information such that at each history* h*, the following hold:*

(a) At least $|A(h)| - 1$ *actions at* h *are clinching actions, and the remaining action (if there is one) is a passing action.*

(b) For every payoff $x \in G_i(h)$*, there exists an action* $a_x \in A(h)$ *that clinches* x *and, following action* a_x*, all of* i*'s remaining moves are trivial moves at which there is only one possible action.*

Proof. For any history h, Lemma 3 shows that for all but at most one action a, all payoffs that are possible following a are also guaranteeable at h. However, it may still be that some action a can lead to multiple payoffs, where different payoffs are guaranteeable for i by following different actions in the future of the game. We now construct a new obviously strategyproof game Γ' that is equivalent to Γ and satisfies the claim of the lemma.

Consider some history h at which i moves. By Lemma 3, all but at most one action (a^*) at h satisfy $P_i((h,a)) \subseteq G_i(h)$; this means that any greedy strategy followed by i with type $\succ_i$ that does not choose a^* guarantees the $\succ_i$-best possible outcome for i. Thus the set $\mathcal{S}_i(h) = \{S_i : y \sim_i x \text{ for all } x, y \in X_i(h, S_i)\}$ contains all possible greedy strategies of agent i for which h is on path and that don't choose a^*. Notice that $\mathcal{S}_i(h)$ is the set of strategies that guarantee some payoff x for i if i plays strategy S_i starting from history h.

We create a new game Γ' that is the same as Γ, except we replace the subgame starting from history h with a new subgame defined as follows. If there is a passing action at h in the original game, then there is a passing action at h in the new game, and the subgame following the passing action is exactly the same as in the original game Γ. Additionally, there are $M = |\mathcal{S}_i(h)|$ other actions at h, denoted $a_1, \dots, a_M$. Each a_m corresponds to one strategy $S_i^m \in \mathcal{S}_i(h)$, and following each a_m, we replicate the original game following the corresponding action $a = S_i^m(h)$ in the original game, except that at any future history $h' \supseteq h$ at which i is called on to act, all actions (and their subgames) are deleted except for the action $a' = S_i^m(h')$ that she would have played in the original game had she followed strategy $S_i^m(\cdot)$. In other words, if i were to choose some action a other than the passing action at h in the original game, then, in the new game Γ', we ask agent i to choose not only her current action, but all future actions as well. By doing so, we have created a new game in which every action (except for the passing action, if it exists) at h clinches some payoff x, and further, agent i is never called upon to move again.[42]

[42] More precisely, all of i's future moves are trivial moves in which she has only one possible action; hence these moves may be removed. Note that this only applies to the non-passing actions at h, which correspond to the guaranteeing strategies in $\mathcal{S}_i(h)$. It is still possible for i to follow the passing action at h and be called upon to make a non-trivial move again later in the game.

We define Γ' counterparts of strategies from Γ so that for all agents $j \neq i$, they continue to follow the same action at every history as they did in the original game, and for i, at history h in the new game, she takes the action a_m that is associated with the strategy S_i^m in the original game. By definition if all the agents follow strategies in the new game analogous to the their strategies from the original game, the same terminal history will be reached, and so Γ and Γ' are outcome-equivalent.

We must also show that if a strategy profile is obviously dominant for Γ, this modified strategy profile is obviously dominant for Γ'. The modified strategy profile is obviously dominant for i because if her obviously dominant action in the original game was to guarantee some payoff x, she now is able to clinch x immediately, which is clearly obviously dominant; if her obviously dominant strategy was to pass at h, she is still able to pass, and the passing action obviously dominates any of the clinching actions (because passing in the original game obviously dominated a strategy that guaranteed any payoff in $G_i(h)$). In addition, the game is also obviously strategyproof for all $j \neq i$ because, prior to h, the set of possible payoffs for j is unchanged, while for any history succeeding h where j is to move, having i make all of her choices earlier in the game only shrinks the set of possible outcomes for j, in the set inclusion sense. When the set of possible outcomes shrinks, the best possible payoff from any given strategy only decreases (according to j's preferences) and the worst possible payoff only increases, and so, if a strategy was obviously dominant in the original game, it will continue to be so in the new game. Repeating this process for every history h, we are left with a new game where, at each history, there are only clinching actions plus (possibly) one passing action, and, following any clinching action, an agent never acts again. □

Step 3: If agent i passes at a history h, then the payoff she ultimately receives must be at least as good as any of the payoffs she could have clinched at h.

This step provides the key insight of why we can have a passing action while still retaining obvious strategy-proofness: whenever an agent passes at a history h, she knows that she will do no worse in the future than anything she could have clinched today.

Lemma 5. *(Necessity) If a game Γ is OSP, then it is equivalent to a millipede game.*

Proof. We have already shown that any OSP game is equivalent to a game with a clinch-pass structure, and the only remaining part is to prove that at least one of conditions (a)-(d) holds. Let history h^i and payoff z be such that (a), (b) and (c) do not hold at h^i, i.e., the following are true:

(a') $z \notin P_i(h^i)$

(b') $z \in P_i(h)$ for all $h \in \mathcal{H}_i(h^i)$

(c') $z \notin \cup_{\tilde{h} \in \mathcal{H}_i(h^i)} C_i(\tilde{h})$ and

Points (b') and (c') imply that z is possible at every $h \subsetneq h^i$ where i is to move, but it is not clinchable at any of them. This implies that for any type of agent i that ranks z first, any obviously dominant strategy must have the agent passing at all $h \in \mathcal{H}_i(h^i)$.[43]

Towards a contradiction, assume that (d) did not hold, i.e., there exists some $h' \in \mathcal{H}_i(h^i)$ and $x \in C_i(h')$ such that $x \notin C_i(h^i)$. Consider a type $z \succ_i x \succ_i \cdots$. We argue that if (d) does not hold at h^i, then there is some $\hat{h}^i \subsetneq h^i$ such that type $\succ_i$ has no obviously dominant action. First, note that at any such $\hat{h}^i \subsetneq h^i$, no clinching action can be obviously dominant, because z is always possible following the passing action, but is never clinchable, and so the worst case from clinching is strictly worse than the best case from passing, which is z. Next, there must be some $\hat{h}^i \subsetneq h^i$ such that the passing action also is not obviously dominant. To see why, note that h^i must be on the path of play for type $\succ_i$, since she must pass at all $h' \subsetneq h^i$. By assumption, $z \notin P_i(h^i)$ and $x \notin C_i(h^i)$, which implies that the worst case outcome from passing at any $h' \subsetneq h^i$ is strictly worse than x. However, we also have $x \in C_i(\hat{h}^i)$ for some $\hat{h}^i \subsetneq h^i$, and so, the best case outcome from clinching x at $\hat{h}^i$ is x. This implies that passing is not obviously dominant, which contradicts that Γ is OSP. □

This last lemma shows that OSP games are equivalent to millipede games. It remains to show that every millipede game is OSP.

Lemma 6. *(Sufficiency) Millipede games are OSP.*

Proof. By definition of a millipede game, there is only one passing action at each node, and so the greedy strategy is essentially unique.[44] For any subset of outcomes $X' \subset \mathcal{X}$, define $Top(\succ_i, X')$ as the best possible payoff in the set X' according to preferences $\succ_i$, i.e., $x \in Top(\succ_i, X')$ if and only if $x \succsim_i y$ for all $y \in X'$ (note that we use our standard convention of identifying any specific outcome x with its indifference class, and so $Top(\succ_i, X')$ is effectively unique). Then, $Top(\succ_i, P_i(h))$ denotes i's top payoff among all payoffs that are possible at history h, and $Top(\succ_i, C_i(h))$ denotes i's top payoff among all of his clinchable payoffs at h.

We argue that if a game Γ is a millipede game, then for each type $\succ_i$ of each agent i, following a greedy strategy is obviously dominant. It is clear that if $Top(\succ_i, C_i(h)) =$

[43] At all such h, since z is not clinchable, but is possible, it must be possible following the (unique) passing action. Thus, no clinching action can be obviously dominant, because the best case outcome from passing is z, while the worst case outcome from clinching is strictly worse than z.

[44] In other words, when playing a greedy strategy, agent i takes a clinching action as soon as her top remaining payoff is clinchable. While for any agent i, there may be multiple ways for her to clinch some payoff x at any given history h, from i's perspective, after each clinching action, i never moves again, and all final histories are payoff-equivalent.

$Top(\succ_i, P_i(h))$, then the greedy strategy of clinching the top payoff is obviously dominant h.[45] What remains to be shown is if $Top(\succ_i, C_i(h)) \neq Top(\succ_i, P_i(h))$, then passing is obviously dominant at h.

Let h be a history such that $Top(\succ_i, C_i(h)) \neq Top(\succ_i, P_i(h))$. To shorten notation, let $x_P(h) = Top(\succ_i, P_i(h))$, $x_C(h) = Top(\succ_i, C_i(h))$, and let $x_W(h)$ be the worst possible outcome from passing at h. Let h be an earliest history such that for all $h' \in \mathcal{H}_i(h)$, we have $x_W(h') \succsim_i x_C(h')$, but $x_W(h) \not\succsim_i x_C(h)$; that is h is an earliest history for which passing is not obviously dominant.

First, note that $h' \in \mathcal{H}_i(h)$ implies $x_W(h) \succsim_i x_W(h')$. Since passing is obviously dominant at all $h' \in \mathcal{H}_i(h)$, we have $x_W(h') \succsim_i x_C(h')$, and together, these imply that $x_W(h) \succsim_i x_C(h')$ for all $h' \in \mathcal{H}_i(h)$. At h, since passing is not obviously dominant, we have $x_C(h) \succ_i x_W(h)$, and further, there must be some $x' \in P_i(h) \setminus G_i(h)$ such that $x' \succ_i x_C(h) \succ_i x_W(h)$ (there may be multiple such x'). The above implies that $x' \succ_i x_C(h) \succ_i x_C(h')$ for all $h' \in \mathcal{H}_i(h)$. Let $X_0 = \{x' \in \mathcal{X} : x' \in P_i(h) \text{ and } x' \succ_i x_C(h)\}$. In words, X_0 is a set of payoffs that are possible at all $h' \subseteq h$, and are strictly better than anything that was clinchable at any $h' \subseteq h$ (and therefore have never been clinchable themselves). Order the elements in X_0 according to $\succ_i$, and wlog, let $x_1 \succ_i x_2 \succ_i \cdots \succ_i x_M$.

We next use the following definition.

Definition. A payoff x **becomes impossible** for i at h' if: (i) $x \in P_i(h'')$ for all $h'' \in \mathcal{H}_i(h')$ and (ii) $x \notin P_i(h')$.

In other words, the history h' at which a payoff becomes impossible is the earliest history where x "disappears" as a possible outcome for i.

Consider a path of play starting from h and ending in a terminal history $\bar{h}$ at which type $\succ_i$ of agent i receives his worst case outcome $x_W(h)$ by following a greedy strategy. For every $x_m \in X_0$, let h_m denote the history on this path at which x_m becomes impossible. Note that because i is receiving $x_W(h)$, such a history h_m exists for all $x_m \in X_0$. Let $\hat{h} = \max\{h_1, h_2, \ldots, h_M\}$ (ordered by $\subset$); in words, $\hat{h}$ is the earliest history at which everything in X_0 is no longer possible.

Claim 1. For all $x_m \in X_0$, payoff x_m is not clinchable for i at any $h_m \subseteq \bar{h}$.

Proof. First, note that x_m is not clinchable at any $h' \subseteq h$ by construction. We will show that x_m cannot be clinchable at any $h' \supset h$ as well. Recall $\hat{h}$ is the earliest history on the path from h to $\bar{h}$ such that all elements of X_0 are no longer possible for i. Further, let let $\hat{h}_{-m} = \max\{h_1, \ldots, h_{m-1}\}$, i.e., $\hat{h}_{-m}$ is the earliest history at which all payoffs strictly

[45]Note that this payoff may also be possible/guaranteeable following the passing action, but since we only need to exhibit one set of obviously dominant strategies, this is sufficient.

preferred to x_m are no longer possible. Assume that x_m were clinchable at some history h' on the path.

Start by considering $m = 1$. At h_1, $x_1 \in P_i(h'')$ for all $h'' \in \mathcal{H}_i(h_1)$, and $h'' \notin P_i(h_1)$, so neither (a) nor (b) hold for x_1 at h_1. Further, we showed above that $x_1 \succ_i x_C(h'')$ for all $h'' \in \mathcal{H}_i(h_1)$, and so x_1 has never been clinchable, and (c) does not hold either. By condition (d) in the definition of a millipede game, at h_1, everything that was ever clinchable before must also be clinchable at h_1, including y. If $X_0 = \{x_1\}$, then at h_1, y is the best possible payoff remaining, and so, greedy strategies require that agent i take y at h_1, which contradicts that she receives $x_W(h)$.

Now, consider an arbitrary m, and assume that for all $m' = 1, \ldots, m-1$, payoff $x_{m'}$ is not clinchable at any h', but x_m is clinchable at some $h_m \subseteq \bar{h}$. Let $x_{m'}$ be (a) payoff that becomes impossible at $\hat{h}_{-m}$. There are two cases:

Case (i): $h_m \subset \hat{h}_{-m}$**.** This is the case where x_m becomes impossible before $x_{m'} \succ_i x_m$. However, note then at history $\hat{h}_{-m}$, neither (a) nor (b) hold for $x_{m'}$. Further, by the inductive hypothesis, $x_{m'}$ is nowhere clinchable, and so (c) does not hold either. This implies that (d) must hold, and so x_m must be clinchable at $\hat{h}_{-m}$. Then, since all preferred payoffs are no longer possible at $\hat{h}_{-m}$, x_m is the best possible payoff remaining, and is clinchable. Therefore, greedy strategies instruct agent i to clinch x_m, which contradicts that she receives $x_W(h)$.

Case (ii): $h_m \supseteq \hat{h}_{-m}$**.** In this case, x_m becomes clinchable after all strictly preferred payoffs are no longer possible. Thus, again, greedy strategies instruct i to clinch x_m, which contradicts that she is receiving $x_W(h)$. □

To finish the proof, again let $\hat{h} = \max\{h_1, h_2, \ldots, h_M\}$ and let $\hat{x}$ be a payoff that becomes impossible at $\hat{h}$. The claim shows that $\hat{x}$ is not clinchable at any $h' \subseteq \hat{h}$, and so condition (d) implies that y must be clinchable at $\hat{h}$. However, if y is clinchable, then greedy strategies direct i to clinch y, which contradicts that she receives $x_W(h)$. This contradiction shows that the game is OSP. □

B Proof and Extensions of Theorem 1

Let us initially assume that each object has one copy, and each agent has single-unit demand. We relax these assumptions below.

Step 1. Let ϕ by an OSP, PE, and symmetric mechanism. By Theorem 2, there is an equivalent millipede mechanism. Since our proof of Theorem 2 gives us a constructive way to derive the millipede representation of ϕ, and the construction does not depend on agents' names, we may assume that this millipede is also symmetric. By definition mechanisms are pruned in the sense of Li, and so the millipede is also pruned, and a mechanism. The greedy strategy is thus the unique obviously dominant strategy.

Step 2. Notice that the symmetric millipede mechanism is a randomization over component mechanisms such that each one is a uniform randomization over a deterministic millipede game. It is enough to show that each such component mechanism is equivalent to RSD.

Step 3. Notice that we can equivalently describe the component mechanism so that nature at its initial move only selects the moving agent and his initial choice set, and then nature moves after every move by an agent, and selects the next mover and the next choice set.[46] By symmetry, we can assume that nature chooses the next mover uniformly at random from among the eligible agents (that is either picks a particular agent who moved already, or picks a random agent who has not moved yet), and that the agent's choice set does not depend on the agent drawn but only on the history of past moves. Finally, let us represent this simple component mechanism in terms of nature initially uniformly randomizing over deterministic millipede games such that only one among the agents who didn't move yet can be the next over.

To define uniform randomizations, for any mechanism ψ and permutation $\sigma : N \to N$ let us denote $\psi^{\sigma}(\succ)(i, a) = \psi\left(\succ_{(\sigma(1),...,\sigma(|N|))}\right)(\sigma(i), a)$. A mechanism $\phi : \mathcal{P}^N \to \mathcal{M}$ is *a uniform randomization* over $\psi : \mathcal{P}^N \to \mathcal{M}$ if

$$\phi\left(\succ_{(1,...,|N|)}\right)(i, a) = \sum_{\sigma:N \overset{1-1}{\to} N} \frac{1}{|N|!} \psi^{\sigma}(\succ)(i, a)$$

For instance, if ψ is a serial dictatorship then ϕ is Random Priority.

Step 4. Let us now fix a preference profile, and construct a bijection from the deterministic millipede games of Step 3 to serial dictatorships. Let's thus take a deterministic millipede game ψ and construct the ordering of agents for the corresponding serial dictatorship. Let i be the first agent who moves along the path of the game, and x the object he receives. If x is the top choice for i, then let's make i the first agent in the ordering. Otherwise, let $x_{i:1}, ..., x_{i:k_i}$ be the unassigned objects that this agent prefers over x while clinching x. By efficiency, these objects are assigned to other agents. Let $i_{i:1}$ be the agent who is assigned x_1. If x_1 is this agent's top choice, then let's put $i_{i:1}$ at the top of the list; otherwise let $x_{i_{i:1}:1}, ..., x_{i_{i:1}:k_{i_{i:1}}}$ be the objects agent $i_{i:1}$ prefers over $x_{i:1}$. Proceeding in this way, we find

[46] In fact, any millipede mechanism can be equivalently described so that nature at its initial move only selects the moving agent and his initial choice set, and then nature moves after every move by an agent, and selects the next mover and the next choice set

the first agent on the list. Say it was $i_{i:1}$. We then look at $i_{i:2}$ and repeat this procedure in an analogous way; if an agent is already in the ordering we skip this agent. Proceeding in this way we construct a list of agents all the way till i such that if they move in order in a serial dictatorship, then each one of them obtains the best still available object. We then look at the first agent i^2 not yet in the order who moved and was assigned object x^2, and analogously as before we add additional agents to the ordering.

The above mapping is well-defined, and running the resulting serial dictatorship for the fixed preference profile results in the same outcome as running the millipede game, because in the millipede game the agent whose object is clinched obtains his most preferred object among objects that were not clinched earlier. (To be sure: the two mechanisms might differ when run on other preference profiles). This mapping is also injective (i.e. one-to-one). Furthermore, because the simple component game uniformly randomizes over $|I|!$ deterministic millipede games, and there are also $|I|!$ serial dictatorship, this one-to-one mapping would then be onto, and hence a bijection.

Step 5. Given Step 4, for any fixed preference profile the outcome of the simple component millipede game is the same as the outcome of uniformly randomizing over serial dictatorships, that is the same as the outcome of random priority. QED

Let us now extend the proof to multiple copies and multi-unit demands. Above, we proved Theorem 1 under additional simplifying assumptions that each object has only one copy and each agent demands at most one object. The proof in the general case follows the same steps, with the following differences. To accommodate multi-object (and multi-unit) demands, it is sufficient to change the terminology and substitute agents' allocations $q \in Q$ for objects. To additionally accommodate multiple-copies supply, we need to take additional care. With multiple copies it is possible that after the first several agents moved, irrespective of their choices, the next agent to move still is able to choose any allocation in Q. In such a case, the probabilities the relevant agents are selected for first, or second (etc.) move need not be the same; however, we obtain an equivalent mechanism by making the probabilities equal for all agents who did not move yet. On the other hand, if after some agents moved, the next agent to move might be unable to choose at least one allocation in Q, then the same argument as in Step 3 of the baseline proof suffices to say that at such a history, all agents who are yet to move have equal chances of having the move. QED

Remark. As mentioned in footnote **??**, the insight of Theorem 1 remains true in some interesting environments beyond its assumptions. Suppose, for instance, that $Q = \left\{q \in \times_{o\in\mathcal{O}} \{0, 1, ..., |o|\} \mid \sum\right.$
All the steps of the proof go through unchanged, except that in Step 2 we proceed as follows. Efficiency implies that each unassigned object is in the set of objects available to the agent who moves first following history h. Indeed, by way of contradiction, suppose that with

positive probability i moves first after h and object o isn't in the set of objects available. By efficiency, all copies of object o need to be allocated to other agents; let j be one of these agents. Let o' be an object in agent i's choice set. Consider the preference profile in which agent i most prefers object o and o' is his second-most preferred object, agents who moved already most prefer the objects they obtained, and other agents most prefer object is o' and o is their second-most preferred object. Then i chooses o', and j obtains o. But, then, the resulting allocation is not efficient; a contradiction that proves the claim of this step, and hence the analogue of Theorem 1 for the environment of footnote **??**.

Assuming Strong Obvious Strategy-Proofness allows us to relax the symmetry assumption in Theorem 1 to the equal treatment of equals. An allocation mechanisms satisfies *equal treatment of equals* if whenever two agents have the same preference ranking, then they obtain the same distribution over outcomes.

Theorem 5. *A game is strongly obviously strategy-proof, efficient, and satisfies equal treatment of equals if and only if it is equivalent to Random Priority.*

Proof. For simplicity we prove the result for the classical house allocation problem in which each object has exactly one copy, and each agent demands at most one object; the general case is similar. Suppose a mechanism ϕ is strongly obviously strategy-proof, efficient, and satisfied the equal treatment of equals. Our characterization of SOSP and efficient mechanisms tells us that we can assume that ϕ can be run as follows: at each history, including the empty history, nature chooses an agent from among the agents who did not move yet, and this agent moves. If there are three or more objects or exactly one object still unallocated, then this agent selects his most preferred still available object and sends an additional message. If, for the first time, there are exactly two unallocated objects, then the agent who moves either (i) selects his most preferred object and sends a message, or (ii) has a choice of clinching one of the two objects (and sending a message) or passing. In the latter case, another agent is then selected by nature, chooses his best object, and the agent who passed obtains the remaining object.

The reminder of the proof is by induction on the number k of agents that already moved. Suppose that agents' moves before the k-th move followed the random priority pattern; this is trivially satisfied when $k = 1$. Consider a history h and name agents so that along h agent 1 moved first and chose object o_1, then agent 2 moved and chose object o_2, etc. till agent $k-1$ who chose object o_{k-1}. Suppose there are at least k agents (otherwise the induction is completed) and consider the k-th agent's move.

Suppose first that there are at least three objects left. Each agent who has not moved yet has equal chances to move first after history h. If h is the empty history, then the claim follows immediately from the equal treatment of equals when the agents rank objects in the same way. If only one agent i_1 moved in h and o_1 is the object he chose, then consider object $o \neq o_1$ and the preference profile in which i_1 ranks o_1 first and o second, while all other agents rank objects in the same way and rank o first. By the equal treatment, the latter agents obtain o with identical probabilities; this probability is the sum of being drawn the first to move and being drawn to move after history h. The inductive assumption implies that the first summand is identical for all these agents, and hence the probability each of them is drawn at h is the same. In general, suppose along h agents i_1, i_2, ..., i_{k-1} moved, in this order, and chose objects o_1, o_2,..., o_{k-1}, respectively. Consider object $o \neq o_1, ..., o_{k-1}$ and the preference profile in which each agent i_ℓ, for $\ell = 1, ..., k-1$, ranks objects so that

$$o_1 \succ_{i_\ell} o_2 \succ_{i_\ell} ... \succ_{i_\ell} o_\ell \succ_{i_\ell} o$$

and other objects (and having no object) are ranked below o; all other agents rank objects in the same way and rank object o first. One of these other agents obtains o if and only if he moves after a subhistory of h. By the inductive assumption the probability of obtaining o is the same for this agents after each proper subhistory of h, and by symmetry the total probability is the same too; hence the probability these agents obtain o after history h—that is move after h—is also the same, as was to be shown.

Suppose now that there are exactly two objects o and o' left at history h. First, note that with two objects, Random Priority is the only strategy-proof and efficient mechanism that satisfies equal treatment of equals. Indeed, with two objects efficient mechanisms are ordinarily efficient, in the sense of Bogomolnaia and Moulin (2001), and the claim follows from their work.[47]

Proceeding similarly to our argument above, suppose that agents i_1, i_2, ..., i_{k-1} moved along h, in this order, and chose objects o_1, o_2,..., o_{k-1}, respectively. Consider object $o \neq o_1, ..., o_{k-1}$ and the preference profile in which each agent i_ℓ, for $\ell = 1, ..., k-1$, ranks objects so that

$$o_1 \succ_{i_\ell} o_2 \succ_{i_\ell} ... \succ_{i_\ell} o_\ell \succ_{i_\ell} o$$

and other objects are ranked below o; all other agents rank objects in the same way and rank object o first. One of these other agents obtains o if and only if he moves after h or its subhistory. By the inductive assumption the probability of obtaining o is the same for

[47] They prove that Random Priority is the only strategy-proof and ordinarily efficient mechanism that satisfies equal treatment of equals when there are three objects and three agents, and with two objects Pareto efficiency and ordinal efficiency become equivalent. The two object is much simpler than the three object problem Bogomolnaia and Moulin study, and one can easily verify the above claim without reliance on Bogomlnaia and Moulin's seminal analysis.

this agents after each proper subhistory of h, and by symmetry the total probability is the same too; hence the probability these agents obtain o after history h—that is move after h—is also the same, as was to be shown. Thus the continuation game following h satisfies strategy-proofness, efficiency, and equal treatment of equals, and by the argument above it is equivalent to Random Priority.

C Proof of Theorem 4

By our previous Lemma 4, any OSP game is equivalent to one such that there is at most one non-clinching move at each history. We show that in fact, for every history that is not penultimate to a terminal history, all moves must be clinching moves. By strengthening OSP to strong OSP, following any move, we need only consider the entire set of possible outcomes for i following any action.[48]

We proceed by induction. Consider $N = 2$, and denote $\mathcal{N} = \{i, j\}$ and $\mathcal{X} = \{x, y\}$. Consider any game efficient and SOSP game Γ. Without loss of generality, let the first mover be i, and note that by efficiency, both x and y must be possible for her. Again without loss of generality, assume she can clinch x at the first move (she must be able to clinch at least one of x or y, since there can be at most one non-clinching move). Consider the first agent who is offered the opportunity to clinch y (following starting the game by a series of passes). If this agent is i, then it is equivalent to offer her the opportunity to clinch y at her first move, and the mechanism is again a serial dictatorship. If the first person to be able to clinch y is j, then it is equivalent to offer her the opportunity to clinch y at her first move, and the game is an almost sequential dictatorship.

Consider now $N = 3$, where $\mathcal{N} = \{i, j, k\}$ and $\mathcal{X} = \{x, y, z\}$, and let the first mover be i. By efficiency, all items are possible for her at the initial history. Assume she had a non-clinching move. This means for one of her actions, labeled a^*, there are (at least) two possible outcomes, x and y, at least one of which (say x) is not clinchable at the initial history. There are two cases, depending on whether the third outcome z is clinchable or not:

z is clinchable at the initial history: By assumption, z is clinchable and x is not. Consider type $x \succ_i z \succ_i y$. None of her clinching actions are strongly obviously dominant, since x is possible following a^*. In addition, a^* is is also not strongly obviously dominant, since y is possible, but she could have clinched z. Thus, this type of agent i has no strongly obviously dominant strategy.

[48] This is slightly subtle under OSP, because the current mover may have "veto power" over some future outcomes, but not others; however, this requires reasoning about the future, and so is eliminated by strong OSP.

z is not clinchable at the initial history: In this case, y is clinchable, but x and z are not (since only one passing move is allowed). Then, consider type $z \succ_i y \succ_i x$. z must be possible (by efficiency), and so must be possible following a^*. This means that no clinching action is strongly obviously dominant. Following the (unique) non-clinching action is also not strongly obviously dominant, because x is possible following a^*, while y is clinchable.

Thus, the first agent to move must have only clinching actions, and, by efficiency, must be able to clinch any object. Following any such clinching move, the game is equivalent to a game of size $N = 2$, which we have already shown is equivalent to an almost-sequential dictatorship.

Last, assume that for every market of size $n = 1, \dots, N-1$ any efficient and SOSP game is equivalent to an almost sequential dictatorship. Consider a market of size N. Let $\mathcal{N} = \{i_1, \dots, i_N\}$ and $\mathcal{X} = \{x_1, \dots, x_N\}$. By efficiency, all items are possible for the first mover, i_1, at the initial history. We argue that all of her actions must be clinching actions.

Assume not. Then there is exactly one action a^* that is a passing action. By definition of a passing action, there must be (at least) two possible outcomes, x_1 and x_2, at least one of which (x_1, say) is not clinchable at the initial history. There are two cases:

There exists a $z \neq x_1, x_2$ that is clinchable at the initial history: By assumption, z is clinchable and x_1 is not. Consider type $x_1 \succ_{i_1} z \succ_{i_1} \cdots$. None of the clinching actions are strongly obviously dominant, since x_1 is possible following a^*, but cannot be clinched. In addition, a^* is not strongly obviously dominant, because x_2 is possible following a^*, while z is clinchable.

There does not exist a $z \neq x_1, x_2$ that is clinchable at the initial history: In this case, x_2 must be clinchable, while all $z \neq x_2$ are not. Choose some $z \neq x_1, x_2$, and consider the type $z \succ_{i_1} x_2 \succ_{i_1} x_1$. Since z is possible (by efficiency), clinching x_2 is not strongly obviously dominant. However, since x_1 is possible following a^*, while x_2 is clinchable, a^* is not strongly obviously dominant either.

Thus, we have that the first mover, i_1, must have only clinching actions, and she must be able to clinch everything. Following any of i_1's clinching actions, we have a game of size $N - 1$, which, by the inductive hypothesis, is an almost sequential dictatorship. It is then simple to see that the overall game is also an almost sequential dictatorship.

D Extensions: Outside Options

Consider the allocation model of Section 3, and suppose that each agent has an outside option.

D.1 Individual Rationality

We say that a game is individually rational if each agent can obtain at least his outside option. The analogues of our results hold true for individually rational games as soon as the domain of each agent's preferences satisfy the domain condition from Sections 2 and 3 restricted to sets $X \subseteq \mathcal{X}$ that do not contain the outside option of this agent. Our proofs remain valid in this setting.

D.2 Restricted Domains

Our results hold true also in allocation domains in which any agent prefers any object to the outside option. Our proofs remain valid also in this setting.[49]

References

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

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

Arribillaga, R. P., J. Massó, and A. Neme (2016): "Not All Majority-based Social Choice Functions Are Obviously Strategy-proof," .

Arrow, K. J. (1963): *Social Choice and Individual Values*, New York: Wiley, 2nd edition ed.

Ashlagi, I. and Y. A. Gonczarowski (2016): "No stable matching mechanism is obviously strategy-proof." Working Paper.

Bade, S. and Y. Gonczarowski (2016): "Gibbard-Satterthwaite Success Stories and Obvious Strategyproofness," .

Bogomolnaia, A. and H. Moulin (2001): "A New Solution to the Random Assignment Problem," *Journal of Economic Theory*, 100, 295–328.

Budish, E. and E. Cantillon (2010): "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard," University of Chicago and Universite Libre de Bruxelles, unpublished mimeo.

Carroll, G. (2010): "A General Equivalence Theorem for Allocation of Indivisible Objects," MIT, Working paper.

Che, Y.-K. and F. Kojima (2010): "Asymptotic Equivalence of Random Priority and Probabilistic Serial Mechanisms," *Econometrica*, 78, 1625–1672.

[49]Note that in this setting every game is individually rational, hence this observation is contained in the previous one.

Dur, U. M. and M. U. Ünver (2015): "Two-Sided Matching via Balanced Exchange: Tuition and Worker Exchanges," NC State University and Boston College, Working Paper.

Ehlers, L. (2002): "Coalitional Strategy-Proof House Allocation," *Journal of Economic Theory*, 105, 298–317.

Gibbard, A. (1973): "Manipulation of Voting Schemes: A General Result," *Econometrica*, 41, 587–601.

——— (1977): "Manipulation of Schemes That Mix Voting with Chance," *Econometrica*, 45, 665–681.

Hakimov, R. and O. Kesten (2014): "The Equitable Top Trading Cycles Mechanism for School Choice," WZB and Carnegie Mellon University, Working Paper.

Hatfield, J. W. (2009): "Strategy-Proof, Efficient, and Nonbossy Quota Allocations," *Social Choice and Welfare*, 33 No. 3, 505–515.

Jones, W. and A. Teytelboym (2016): "Choices, preferences and priorities in a matching system for refugees," *Forced Migration Review*.

Li, S. (2016): "Obviously Strategy-Proof Mechanisms," Working Paper.

Liu, Q. and M. Pycia (2013): "Ordinal Efficiency, Fairness, and Incentives in Large Markets," Working Paper.

Loertscher, S. and L. M. Marx (2015): "Prior-Free Bayesian Optimal Double-Clock Auctions," Working Paper.

Morrill, T. (2014): "Making Just School Assignments," *Games and Economic Behavior*, (forthcoming).

Pápai, S. (2000): "Strategyproof Assignment by Hierarchical Exchange," *Econometrica*, 68, 1403–1433.

——— (2001): "Strategyproof and Nonbossy Multiple Assignments," *Journal of Public Economic Theory*, 3, 257–271.

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

Pycia, M. (2011): "Ordinal Efficiency, Fairness, and Incentives in Large Multi-Unit-Demand Assignments," .

Pycia, M. and M. U. Ünver (2009): "Incentive Compatible Allocation and Exchange of Discrete Resources," UCLA and Boston College, unpublished mimeo.

——— (2016): "Arrovian Efficiency in Allocation of Discrete Resources," UCLA and Boston College, Working Paper.

Rosenthal, R. W. (1981): "Games of perfect information, predatory pricing and the chain-store paradox," *Journal of Economic theory*, 25, 92–100.

Roth, A. (2015): “Migrants Aren’t Widgets,” *Politico Europe*, http://www.politico.eu/article/migrants–arent–widgets–europe–eu–migrant–refugee–crisis/.

Roth, A. E., T. Sönmez, and M. U. Ünver (2004): “Kidney Exchange,” *Quarterly Journal of Economics*, 119, 457–488.

Satterthwaite, M. (1975): “Strategy-proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions,” *Journal of Economic Theory*, 10, 187–216.

Shapley, L. and H. Scarf (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 23–37.

Sönmez, T. and M. U. Ünver (2010): “Course Bidding at Business Schools,” *International Economic Review*, 51, 99–123.

Troyan, P. (2016): “Obviously Strategyproof Implementation of Allocation Mechanisms,” Working paper, University of Virginia.