

A Theory of Simplicity in Games and Mechanism Design

Marek Pycia and Peter Troyan*

Abstract

We introduce a general class of simplicity standards that vary the foresight abilities required of agents in extensive-form games. Rather than planning for the entire future of a game, agents are presumed to be able to plan only for those histories they view as simple from their current perspective. Agents may update their so-called strategic plan as the game progresses, and, at any point, for the called-for action to be simply dominant, it must lead to unambiguously better outcomes, no matter what occurs at non-simple histories. We use our approach to simplicity to provide characterizations of simple mechanisms in general social choice environments both with and without transfers, including canonical mechanisms such as ascending auctions, posted prices, and serial dictatorship-style mechanisms. As a final application, we explain the widespread popularity of the well-known Random Priority mechanism by characterizing it as the unique mechanism that is efficient, fair, and simple to play.

*Pycia: University of Zurich; Troyan: University of Virginia. First presentation: April 2016. First posted draft: June 2016. We initially presented and distributed this paper under the title “Obvious Dominance and Random Priority.” For their comments, we would like to thank Itai Ashlagi, Eduardo Azevedo, Dirk Bergemann, Ernst Fehr, Ben Golub, Yannai Gonczarowski, Ed Green, Rustamdjan Hakimov, Fuhito Kojima, Shengwu Li, Giorgio Martini, Stephen Morris, Anna Myjak-Pycia, Nick Netzer, Ran Shorrer, Erling Skancke, Utku Ünver, Roberto Weber, anonymous referees, the Eco 514 students at Princeton, and the audiences at the 2016 NBER Market Design workshop, NEGT’16, NC State, ITAM, NSF/CEME Decentralization, the 2017 Econometric Society Meetings, UBC, the Workshop on Game Theory at NUS, UVa, ASSA 2019, MATCH-UP 2019, ESSET 2019, and EC’19, where this paper was awarded the Best Paper prize. Samuel Haefner and Simon Lazarus provided excellent research assistance. Pycia gratefully acknowledges the support of the UCLA Department of Economics and the William S. Dietrich II Economic Theory Center at Princeton. Troyan gratefully acknowledges support from the Bankard Fund for Political Economy and the Roger Sherman Fellowship at the University of Virginia.

1 Introduction

Consider a group of agents who must come together to make a choice from some set of potential outcomes that will affect each of them. This can be modeled as having the agents play a “game”, taking turns choosing from sets of actions (possibly simultaneously), with the final outcome determined by the decisions made by all of the agents each time they were called to play. To ensure that the ultimate decision taken satisfies desirable normative properties (e.g., efficiency), the incentives given to the agents are crucial. The standard route taken in mechanism design is to appeal to the revelation principle and use a Bayesian or dominant-strategy incentive-compatible direct mechanism where agents are simply asked to report their private information, and it is always in their interest to do so truthfully. However, this is useful only to the extent the participants understand that being truthful is in their interest, and indeed, there is evidence many real-world agents do not report the truth, even in strategy-proof mechanisms.¹ In other words, Bayesian or dominant-strategy mechanisms, while theoretically appealing, may not actually be easy for participants to play in practice.

What mechanisms, then, are actually “simple to play”? We address this question for a broad range of social choice environments both with and without transfers. We do so by introducing a general class of simplicity standards that vary the foresight abilities required of agents in extensive-form imperfect-information games, and use them to provide characterizations of simple mechanisms for a wide range of settings; in particular, we provide microfoundations for popular simple mechanisms such as posted prices and Random Priority.

Our general approach relies on the idea that, rather than planning a strategy for any possible future point that may be reached in a game, a player plans for only those nodes (or information sets) that he or she perceives as simple from the current perspective; formally, we refer to such objects as *strategic plans*.² Then, for a strategic plan to be *simply dominant*, the called for action needs to be unambiguously better than alternatives, irrespective of what happens at information sets that are not simple for the agent. As the game progresses, the agent’s perception of which information sets are simple may change. Importantly, we allow for the possibility that the agent can change his or her strategic plan along the path of the game, which differentiates strategic plans from the standard game-theoretic concept of a strategy. The sets of information sets perceived as simple from the perspective of a

¹Kagel, Harstad, and Levin (1987), Li (2017b), Hassidim, Romm, and Shorrer (2016), Rees-Jones (2017), Rees-Jones (2018), Shorrer and Sóvágó (2018), and Artemov, Che, and He (2017).

²Savage (1954) wrestles with whether decision-makers should be modeled as “look before you leap” (create a full strategic plan for all possible future decisions one may face) or “you can cross that bridge when you come to it” (make choices as they arise). While standard strategic concepts of game theory formalize the former modeling option, our approach formalizes the latter.

given information set are taken as a primitive of our definition, and the smaller (in an inclusion sense) are these sets of simple information sets, the stronger is the resulting simplicity standard.³

One important special case covered by our approach is Li’s (2017b) obvious dominance. A strategy in an imperfect-information extensive-form game is *obviously dominant* if, whenever an agent is called to play, even the worst possible final outcome from following the prescribed strategy is at least as good as the best possible outcome from any other strategy, where the best and worst cases are determined by considering all possible strategies that could be played by her opponents in the future, keeping the agent’s own strategy fixed. Our general approach to simplicity captures obviously dominant strategies when agents perceive all of their own information sets as simple and all information sets of other agents as not simple—in other words, at each information set, agents are able to plan the action that they will take at any future information set at which they may be called to play. This highlights an important feature of obvious dominance, which is that it presumes that agents can perform demanding backward induction over at least their own future actions. As an example, consider chess: assuming that White can always force a win, any winning strategy of White is obviously dominant; yet, the strategic choices in chess are far from obvious.

To get a better understanding of what strategies are indeed simple, we also analyze more demanding standards in our class. The first is *one-step-foresight (OSF) dominance*. A strategic plan is *one-step-foresight dominant* if it is simply dominant for players who perceive as simple their current information set and only the first information sets at which they may be called to play in the continuation game; in other words, agents are able to plan at most one move ahead at a time. For instance, in an ascending auction planning to stay in is dominant for such players as long as the current price is below their value because they can foresee the next round of bidding and they can always drop out at the next round (notice that our framework allows for the bidder’s strategic plan to be adjusted when the next round is actually reached). A strategic plan is *strongly obviously dominant* if it is simply dominant for players who perceive as simple only their current information set. In other words, a strategic plan is strongly obviously dominant if, whenever an agent is called to play, even the worst possible final outcome from the prescribed 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. Thus, strongly obviously dominant plans are those that are weakly better than all alternative actions even if the agent is concerned

³We show that a strategic plan is simply dominant if and only if in every game an agent can confuse with the actual game the strategic plan is weakly dominant in the standard sense (Theorem 2). A related behavioral microfoundation for his obvious dominance standard—on which we build—was provided by Li (2017b).

that she might tremble in the future or has time-inconsistent preferences.

For each of these three subclasses of our general simplicity construction, we analyze which games are simple. For obvious dominance, we focus on social choice environments without transfers, hence complementing Li (2017b), who focused on the case with transfers.⁴ We call the class of obviously dominant games in these environments *millipede games*. In a millipede game, each time an agent is called to move, she is presented with some subset of payoff-equivalent outcomes, or more simply *payoffs*, that she can ‘clinch’, after which she leaves the game; she also may be given the opportunity to ‘pass’ and remain in the game, with the potential of being offered better clinching options in the future. If this agent passes, another agent is presented with an analogous choice, etc., until one of them eventually clinches. While some millipede games, such as serial dictatorships, are frequently encountered and are indeed simple to play, others are rarely observed in market-design practice, and their strategy-proofness is not necessarily immediately clear. In particular, similar to chess, some millipede games require agents to look far into the future and to perform potentially complicated backward induction reasoning.

We next study one-step-foresight dominance in environments with and without transfers. We show that in environments including single-unit auctions and binary public good choice, any social choice rule that is implementable in obviously dominant strategies is also implementable in one-step-foresight dominant strategies. In particular, any one-step-foresight simple mechanism is equivalent to a personal clock auction as defined by Li (2017b).⁵ We also note that personal clock auctions are monotonic in the following sense: each time an agent is called to move, at any next move in the continuation game at which the agent is called again (or terminal history), the agent is able to clinch a payoff that is either at least as good as anything she could have clinched previously, or at least as good as anything that was possible but not clinchable. For instance, in an ascending auction, the only clinchable payoff is that associated with dropping out, except if the agent wins. We further show that in no-transfer environments, all one-step-foresight simple millipedes are similarly monotonic, thereby eliminating the complex—yet still OSP—millipede games identified previously.

We also study strong obvious dominance in environments with and without transfers. We show that strongly obviously strategy-proof games do not require agents to look far into the future and perform lengthy backwards induction: in all such games, each agent

⁴Social choice problems without transfers are ubiquitous in the real-world. Examples include voting (Arrow, 1963), school choice (Abdulkadiroğlu and Sönmez, 2003), organ exchange (Roth, Sönmez, and Ünver, 2004), course allocation (Sönmez and Ünver, 2010; Budish and Cantillon, 2012), and refugee resettlement (Jones and Teytelboym, 2016; Delacrétaz et al., 2016).

⁵Our analysis in environments with transfers builds on Li’s characterization of obviously dominant implementation in these environments.

has essentially at most one payoff-relevant move. Building on this insight, we show that all strongly obvious strategy-proof games can be implemented as sequential price games in which each agent moves at most once, and, at this move, is offered a choice from a menu of options (which may or may not include transfers). If the menu has three or more options for the agent in question, then the agent’s final payoff is what they choose from the menu. If the menu has only two options, then the agent’s final payoff might depend on other agents’ choices, but truthfully indicating the preferred option is the dominant choice. In this way, strong obvious dominance gives us a microfoundation for posted prices, a ubiquitous sales mechanism.⁶

In the final section of the paper, as an application of our approach to simplicity, we provide an axiomatic characterization of the well-known Random Priority (also known as Random Serial Dictatorship) mechanism. In the context of no-transfer allocation problems, Random Priority works as follows: first Nature selects an ordering of agents, and then each agent moves in turn and chooses her favorite object among those that remain available given previous agents’ choices. This mechanism has a long history (cf. Abdulkadiroğlu and Sönmez, 1998) and it is used in a wide variety of practical allocation problems, including school choice, worker assignment, course allocation, and the allocation of public housing. Random Priority is well-known to have good efficiency, fairness, and simplicity properties: it is Pareto efficient, it treats agents in a symmetric way, and it is (obviously) strategy-proof.⁷ However, it has until now remained unknown whether there are other such mechanisms, and if so, what explains the relative popularity of Random Priority over these alternatives.⁸ We show that there are none, thus resolving positively the quest to establish Random Priority as the unique mechanism with good incentive, efficiency, and fairness properties and thereby explaining its popularity in practical market design settings.

Our results build on the key contributions of Li (2017b), who formalized obvious strategy-proofness and established its desirability as an incentive property (see the discussion above). Our construction of the simplicity criteria—while being more general and allowing us to select more precisely simple mechanisms—is inspired by his work. Li’s work generated a substantive

⁶For an earlier microfoundation of posted prices, see Hagerty and Rogerson (1987) and Copic and Ponsati (2016).

⁷For discussion of efficiency and fairness see, e.g., Abdulkadiroğlu and Sönmez (1998), Bogomolnaia and Moulin (2001), Che and Kojima (2010), and Liu and Pycia (2011). Obvious strategy-proofness of Random Priority was established by Li (2017b).

⁸In 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 (2011) and Pycia (2011) proved that Random Priority is the asymptotically unique mechanism that is symmetric, asymptotically strategy-proof, and asymptotically ordinally efficient. While these earlier results looked at either very small or very large markets, ours is the first characterization that holds for any number of agents and objects.

interest focused on his simplicity standard. Following up on Li’s work, but preceding ours, Ashlagi and Gonczarowski (2018) show that stable mechanisms such as Deferred Acceptance are not obviously strategy-proof, except in very restrictive environments where Deferred Acceptance simplifies to an obviously strategy-proof game with a ‘clinch or pass’ structure similar to simple millipede games (though they do not describe it in these terms). Other related papers include Troyan (2019), who studies obviously strategy-proof allocation via the popular Top Trading Cycles (TTC) mechanism, and provides a characterization of the priority structures under which TTC is OSP-implementable.⁹ Following our work, Arribillaga et al. (2017) characterize the voting rules that are obviously strategy-proof on the domain of single-peaked preferences and, in an additional result, in environments with two alternatives; Bade and Gonczarowski (2017) study obviously strategy-proof and efficient social choice rules in several environments. Mackenzie (2017) introduces the notion of a “round table mechanism” for OSP implementation and draws parallels with the standard Myerson-Riley revelation principle for direct mechanisms. There has been less work that goes beyond Li’s obvious dominance. Li (2017a) extends his ideas to an ex post equilibrium context, while Zhang and Levin (2017a; 2017b) provide decision-theoretic foundations for obvious dominance and explore weaker incentive concepts.

More generally, our work also contributes to the understanding of limited foresight and limits on backward induction. Other work in this area—with different approaches from ours—includes Jehiel’s (1995; 2001) studies of limited foresight, Ke’s (2015) axiomatic approach to bounded horizon backward induction, as well as the rich literature on time-inconsistent preferences (e.g., Laibson (1997) and Gul and Pesendorfer (2001; 2004)). The paper also adds to our understanding of dominant incentives, efficiency, and fairness in settings with and without transfers. In settings with transfers, these questions were studied by e.g. Vickrey (1961), Clarke (1971), Groves (1973), Green and Laffont (1977), Holmstrom (1979), Dasgupta et al. (1979), and Hagerty and Rogerson (1987). 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 (2016; 2017) 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.¹⁰

⁹Li showed that the classic top trading cycles (TTC) mechanism of Shapley and Scarf (1974), in which each agent starts by owning exactly one object, is not obviously strategy-proof. Also of note is 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.

¹⁰Pycia 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).

Liu and Pycia (2011), Pycia (2011), Morrill (2014), Hakimov and Kesten (2014), Ehlers and Morrill (2017), and Troyan et al. (2018) characterize mechanisms that satisfy certain incentive, efficiency, and fairness objectives.

2 Model

2.1 Preferences

Let $\mathcal{N} = \{i_1, \dots, i_N\}$ be a set of agents, and \mathcal{X} a finite set of outcomes.¹¹ Each agent has a preference ranking over outcomes, where, for any two $x, y \in \mathcal{X}$, we write $x \succsim_i y$ to denote that x is weakly preferred to y . We allow for indifferences, and write $x \sim_i y$ if $x \succsim_i y$ and $y \succsim_i x$. For any \succsim_i , we let \succ_i denote the corresponding strict preference relation, i.e., $x \succ_i y$ if $x \succsim_i y$ but not $y \succsim_i x$. We will generally work with the strict preference relation \succ_i , which we refer to as an agent’s **type**. The domain of preferences of agent $i \in \mathcal{N}$ is denoted \mathcal{P}_i .

We study both settings with and without transfers. The main assumption we make on the preference domains is that they are rich. Our formalization of richness takes as a primitive a dominance relation over outcomes, denoted \succeq , where \succeq is a reflexive and transitive binary relation on \mathcal{X} . If $x \succeq y$ but not $y \succeq x$, then we write $x \triangleright y$. A preference ranking \succsim_i is **consistent** with \succeq if $x \succeq y$ implies that $x \succsim_i y$ and $x \triangleright y$ implies that $x \succ_i y$. We say that \mathcal{P}_i is **rich** if it consists of all strict rankings that are consistent with \succeq .¹²

We allow that different agents have different preference domains; that is different agents’ preference domains might be governed by different dominance relations, \succeq_i . If $x \succeq_i y$ and $y \succeq_i x$ then x and y are \succeq_i –**equivalent**. Any such \succeq_i determines a partition of \mathcal{X} , which we refer to as an **equivalence partition**. We refer to each element of the equivalence partition as a **payoff** of the agent in question. When the distinction between a payoff and an outcome is important we write $[x]_i = \{y \in \mathcal{X} : x \succeq_i y \text{ and } y \succeq_i x\}$ to represent the payoff (the element of the partition) that contains x . Elsewhere, to avoid unnecessary formalism, we say payoff x to refer to the partition element to which outcome x belongs; thus phrases such as “payoff x obtains” are understood as “some $y \in [x]_i$ obtains”. A payoff x is called **undominated** in a subset of payoffs of agent i if there is no payoff y in this subset such

¹¹Assuming \mathcal{X} is finite simplifies the exposition and is satisfied in the examples 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.

¹²The literature contains many different definitions of rich preference domains, cf. e.g. Dasgupta, Hammond, and Maskin (1979) and Pycia (2012). Our use of the term shares with these earlier uses the idea that the domain of preferences contains sufficiently many profiles; if certain preference profiles belong to the domain of preference profiles then some other profiles belong to the domain as well.

that $y \triangleright_i x$.

Some examples may help in illuminating our modeling framework of dominance and richness, and show how it is very flexible and encompasses many important special cases. Examples without transfers include:

- **Voting:** Every agent has strict preferences over all alternatives in \mathcal{X} . This is captured by the trivial dominance relation \succeq_i in which $x \succeq_i y$ implies $x = y$ for all i . Each agent's preference domain \mathcal{P}_i partitions \mathcal{X} into $|\mathcal{X}|$ individual subsets, and richness implies that every strict preference ranking over \mathcal{X} belongs to \mathcal{P}_i for each i .
- **Allocating indivisible goods without transfers:** Each $x \in \mathcal{X}$ describes the entire allocation of goods to each of the agents. Each agent is indifferent over how goods she does not receive are assigned to others. This is captured by a dominance relation \succeq_i for agent i defined as follows: $x \succeq_i y$ if and only if agent i receives the same set of goods in outcomes x and y . Each element of agent i 's equivalence partition can be identified with the set of objects she receives, and richness implies that every strict ranking of these sets belongs to \mathcal{P}_i for each i .

Formally, whenever the dominance relation \succeq_i is also symmetric for all i (in addition to being reflexive and transitive), we say the environment is a **no-transfer environment**.¹³ Environments with transfers are also covered by our model, though, as transfers put extra structure on the model, they will be governed by dominance relations that are not symmetric. Examples include:

- **Social choice with transfers:** Let $\mathcal{X} = \mathcal{Y} \times \mathcal{W}^N$, where \mathcal{Y} is a set of substantive outcomes, $\mathcal{W} \subsetneq \mathbb{R}$ a (finite) set of possible transfers, and $w \equiv (w_i)_{i \in \mathcal{N}}$ denotes the profile of transfers to the agents. Each agent i prefers to pay less rather than more (for a fixed $y \in \mathcal{Y}$) and is indifferent between any two outcomes that vary only in other agents' transfers. This preference domain is given by the dominance relation $(y, w) \succeq_i (y', w')$ if and only if $y = y'$ and $w_i \geq w'_i$.
- **Binary allocation with transfers.** $\mathcal{Y} \subseteq \{0, 1\}^N$ is a set of feasible allocations and $\mathcal{W} \subsetneq \mathbb{R}$ is a set of transfers, with $\mathcal{X} = \mathcal{Y} \times \mathcal{W}^N$. For any $y \equiv (y_i)_{i \in \mathcal{N}} \in \mathcal{Y}$, $y_i = 1$ denotes that i is in the allocation, and $w \equiv (w_i)_{i \in \mathcal{N}} \in \mathcal{W}^N$ denotes the profile of transfers. The dominance relation for agent i is defined as follows: $(y, w) \succeq_i (y', w')$ if and only if $w_i \geq w'_i$ and $y_i \geq y'_i$. This is the main application studied by Li (2017b), and covers

¹³A binary relation \succeq_i is *symmetric* if $x \succeq_i y$ implies $y \succeq_i x$. It is easy to see that this holds in the examples without transfers above, but not in those with transfers below.

such applications as unit-demand auctions, procurement auctions, and binary public goods problems.

These are just a few examples of settings that fit into our general model. While richness is a very flexible assumption, not all preference domains are rich. For instance, domains of single-peaked preferences are typically not rich.¹⁴

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 $x \in \text{supp}(\mu)$ and $y \in \text{supp}(\nu)$ this agent weakly prefers x over y ; the preference between μ and ν is strict if, additionally, at least one of the preferences between $x \in \text{supp}(\mu)$ and $y \in \text{supp}(\nu)$ 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 .

2.2 Mechanisms

To determine the outcome that will be implemented, the planner designs a game Γ for the agents to play. Formally, we consider imperfect-information, extensive-form games with perfect recall, which are defined in the standard way: there is a finite collection of partially ordered **histories** (sequences of moves), \mathcal{H} . We use the notation $h' \subseteq h$ to denote that $h' \in \mathcal{H}$ is a subhistory of $h \in \mathcal{H}$, and write $h' \subset h$ when $h' \subseteq h$ but $h \neq h'$. Terminal histories (those with no successors) will be denoted with bars, i.e., \bar{h} . Each $\bar{h} \in \mathcal{H}$ is associated with an outcome in \mathcal{X} , and agents receive payoffs at \bar{h} that are consistent with their preferences over outcomes \succ_i . At every non-terminal history $h \in \mathcal{H}$, one agent, denoted i_h , is called to play and has a finite set of **actions** $A(h)$ from which to choose. We write $h' = (h, a)$ to denote the history h' that is reached by starting at history h and following the action $a \in A(h)$. To avoid trivialities, we assume that no agent moves twice in a row and that $|A(h)| > 1$ for all non-terminal $h \in \mathcal{H}$. To capture random mechanisms, we also allow for histories h at which a non-strategic agent, Nature, is called to move, and selects an action in $A(h)$ according to some probability distribution.

The set of histories at which agent i moves is denoted $\mathcal{H}_i = \{h \in \mathcal{H} : i_h = i\}$. The set \mathcal{I}_i is a partition of \mathcal{H}_i into **information sets**, where, for any information set $I \in \mathcal{I}_i$ and $h, h' \in I$ 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

¹⁴We might slightly relax the richness assumption. For instance, in the presence of outside options we would say that a game is individually rational if each agent can obtain at least his outside option. To obtain the analogues of many of our results for individually rational games, it is sufficient to assume that the domain of each agent's preferences satisfies the richness condition restricted to sets $X \subseteq \mathcal{X}$ that do not contain the outside option of this agent. However, Arribillaga, Massó, and Neme (2017) show that some of our results do not extend to single-peaked preference domains.

\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 $I(h) \in \mathcal{I}_i$ the information set containing history h .¹⁵ These imperfect information games allow us to incorporate incomplete information in the standard way in which Nature moves first and determines agents' types. Due to the nature of the dominance properties we study, we do not need to make any assumptions on agents' beliefs about others' types.

A **strategy** for a player i in game Γ is a function S_i that maps i 's information sets into actions chosen by the agent at each information set.¹⁶ When we want to refer to the strategies of different types \succ_i of agent i , we write $S_i(\succ_i)$ for the strategy followed by agent i of type \succ_i ; in particular, $S_i(\succ_i)(I_i)$ denotes the action chosen by agent i with type \succ_i at information set $I_i \in \mathcal{I}_i$. We use $S_{\mathcal{N}}(\succ_{\mathcal{N}}) = (S_i(\succ_i))_{i \in \mathcal{N}}$ to denote the strategy profile for all of the agents when the type profile is $\succ_{\mathcal{N}} = (\succ_i)_{i \in \mathcal{N}}$. An **extensive-form mechanism**, or simply a **mechanism**, is an extensive-form game Γ together with a profile of strategies $S_{\mathcal{N}}$. Two extensive-form mechanisms $(\Gamma, S_{\mathcal{N}})$ and $(\Gamma', S'_{\mathcal{N}})$ are **equivalent** if for every profile of types $\succ_{\mathcal{N}} = (\succ_i)_{i \in \mathcal{N}}$, the distribution over outcomes— $(\Gamma, S_{\mathcal{N}})(\succ_{\mathcal{N}})$ —when agents follow $S_{\mathcal{N}}(\succ_{\mathcal{N}})$ in Γ is the same as when agents follow $S'_{\mathcal{N}}(\succ_{\mathcal{N}})$ in Γ' .¹⁷

3 Example: Obvious Dominance and Millipede Games

How to define the concept of games that are “simple to play”? As an example, we re-examine obvious strategy-proofness, the seminal simplicity standard proposed by Li (2017b). Given a game Γ , a strategy S_i **obviously dominates** another strategy S'_i for player i if, starting from any earliest information set I_i 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 S'_i , where the best/worst case outcomes are determined over all possible strategies of other agents S_{-i} and all possible choices of Nature. A profile of strategies $S_{\mathcal{N}}(\cdot) = (S_i(\cdot))_{i \in \mathcal{N}}$ is **obviously dominant** if for every player i and every type \succ_i , the strategy $S_i(\succ_i)$ obviously dominates every other strategy S'_i . Γ is **obviously strategy-proof (OSP)** if there exists a profile of strategies $S_{\mathcal{N}}(\cdot)$ that is obviously dominant.

While Li (2017b) shows that ascending auctions have obviously dominant strategies (and

¹⁵We will see shortly that it is essentially without loss of generality to assume all information sets are singletons, and so will be able to drop the $I(h)$ notation and identify each information set with the unique sequence of actions (i.e., history) taken to reach it.

¹⁶We consider pure strategies, but the analysis can be easily extended to mixed strategies.

¹⁷The equivalence concept here is outcome-based, and hence different from the procedural equivalence concept of Kohlberg and Mertens (1986).




Figure 1: An example of a millipede game.

second-price sealed bid auctions do not), we show that many obviously dominant strategies are not necessarily simple: for instance, as discussed in the introduction, if White has a winning strategy in chess, then this strategy is obviously dominant. Despite this classification, not only can we not calculate the White’s winning strategy, it is unknown whether such a strategy even exists. This motivates two natural questions: (1) what classes of games admit obviously dominant strategies? and (2) can we formally define a simplicity standard that better delineates classes of games that are generally understood to be simple?

We first tackle the first of these questions and characterize the entire class of obviously strategy-proof mechanisms in **environments without transfers**, which recall from Section 2, are defined by the assumption that \succeq_i is reflexive, transitive, and symmetric for all i . We call the resulting class of games *millipede games*. Roughly speaking, a millipede game is a take-or-pass game similar to a centipede game (Rosenthal, 1981), but with more players and more actions (i.e., “legs”) at each node. Figure 1 shows the extensive form of a millipede game for the special case of object allocation with single-unit demand, where the agents are labeled i, j, k, \dots and the objects are labeled w, x, y, \dots . 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 . If he takes an object, he leaves the game and it 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 1 considers an object allocation environment, we define millipede games more generally for all environments without transfers. Recall that each agent’s preference domain \mathcal{P}_i partitions the outcome space \mathcal{X} into equivalence classes, with each element referred to as a payoff for agent i , and richness in this setting says that every strict ranking of these

payoffs is a possible preference type for i . We say that a payoff x is **possible** for agent i at history h if there is a strategy profile of all the agents (including choices made by Nature) such that, starting from h and following this strategy profile results in a terminal history \bar{h} at which agent i obtains payoff x . For any history h , $P_i(h)$ denotes the set of payoffs that are possible for i at h . We say agent i has **clinched** payoff x at history h if agent i receives payoff x at all terminal histories $\bar{h} \supseteq h$. If after following action $a \in A(h)$, an agent receives the same payoff for every terminal $\bar{h} \supseteq (h, a)$, we say that a is a **clinching** action. If an action $a \in A(h)$ is not a clinching action, then it is called a **passing** action.

We denote the set of payoffs that i can clinch at a history h at which she moves by $C_i(h)$; that is, $x \in C_i(h)$ if there is some action $a \in A(h)$ such that i receives payoff x for all terminal $\bar{h} \supseteq (h, a)$. At a terminal history \bar{h} , no agent is called to move and there are no actions; however, it will be useful to define $C_i(\bar{h}) = \{x\}$ for all i , where x is the outcome that obtains at terminal history \bar{h} . We further define $C_i^{\subseteq}(h) = \{x : x \in C_i(h') \text{ for some } h' \subseteq h \text{ s.t. } i_{h'} = i\}$ to be the set of payoffs that i can clinch at some subhistory of h , and $C_i^{\subset}(h) = \{x : x \in C_i(h') \text{ for some } h' \subsetneq h \text{ s.t. } i_{h'} = i\}$ to be the set of payoffs that i can clinch at some *strict* subhistory of h . Note that the definition of $C_i(h)$ implicitly presumes that $i_h = i$, i.e., i moves at h ; however, $P_i(h)$, $C_i^{\subseteq}(h)$ and $C_i^{\subset}(h)$ are well-defined for any h , whether i moves at h or not.

Finally, consider a history h such that $i_{h'} = i$ for some $h' \subsetneq h$ (i.e., i moves before h), and either $i_h = i$ or h is a terminal history. We say that payoff x **becomes impossible** for i at h if $x \in P_i(h')$ for all $h' \subsetneq h$ such that $i_{h'} = i$, but $x \notin P_i(h)$. We say that payoff x is **previously unclinched** at h if $x \notin C_i^{\subseteq}(h)$.

Given these definitions, we define a **millipede game** as a finite extensive-form game of perfect information that satisfies the following properties:

1. Nature either moves once, at the empty history h_\emptyset , or Nature has no moves.
2. At any other history $h \neq h_\emptyset$, all but at most one action are clinching actions, and the remaining action (if there is one) is a passing action. (Note that there may be several clinching actions associated with the same payoff for the agent who moves at h .)
3. At all h , if there exists a previously unclinched payoff x that becomes impossible for agent i_h at h , then $C_{i_h}^{\subseteq}(h) \subseteq C_{i_h}(h)$.

In a millipede game, if an agent's top still-possible payoff, say x , is not clinched at some history h , it is easy to see that no clinching action can be obviously dominant; the last condition ensures that passing will be obviously dominant, since if x becomes impossible, then the agent will at least be able to return to any payoff she was previously offered to

clinch. Notice that millipede games have a recursive structure: the continuation game that follows any action is also a millipede game. A simple example of a millipede game is a deterministic sequential dictatorship in which no agent has passing moves and all payoffs that are not precluded by the earlier choices of other agents are clinchable. A more complex example is sketched in Figure 1.¹⁸

Our first main result is to characterize the class of OSP games and mechanisms as the class of millipede games with greedy strategies. To define greedy strategies, let $Top(\succ_i, P_i(h))$ denote the best payoff in the set $P_i(h)$ for an agent of type \succ_i . A strategy $S_i(\succ_i)$ is called **greedy** if, for any h at which $Top(\succ_i, P_i(h)) \in C_i(h)$, the action $S_i(\succ_i)(h)$ clinches this payoff for the agent; otherwise, the agent passes.

Theorem 1. *In environments without transfers, every obviously strategy-proof mechanism (Γ, S_N) is equivalent to a millipede game with greedy strategies. Every millipede game with greedy strategies is obviously strategy-proof.*

This theorem is applicable in many environments. This includes allocation problems in which agents care only about the object(s) they receive, in which case, clinching actions correspond to taking a specified object and leaving the remaining objects to be distributed amongst the remaining agents. Theorem 1 also applies to standard social choice problems in which no agent is indifferent between any two outcomes (e.g., voting), in which case clinching corresponds to determining the final outcome for all agents. In such environments, Theorem 1 implies the following:

Corollary 1. *If each agent has strict preferences among all outcomes, then every OSP game is equivalent to a game in which either there are only two outcomes that are possible when the first agent moves (and the first mover can either clinch any of them, or can clinch one of them or pass to a second agent, who is presented with an analogous choice, etc.), or the first agent to move can clinch any possible outcome and has no passing action.*

The latter case of Corollary 1 is the standard dictatorship, with a possibly restricted set of possible outcomes, while the former case is a generalization that allows for the possibility that at her turn, an agent can enforce one of the two outcomes, but not the other (the enforceable option may differ at each agent’s turn). In particular, this corollary gives an analogue of the Gibbard-Satterthwaite dictatorship result, with no efficiency assumption.

The full proof of Theorem 1 is in the appendix; here, we provide a brief sketch of the more interesting direction that for any OSP game Γ , there is an equivalent millipede game.

¹⁸The first more complex example of a millipede game we know of is due to Ashlagi and Gonczarowski (2018). They construct an example of OSP-implementation of deferred acceptance on some restricted preference domains. On these restricted domains, DA reduces to a millipede game (though they do not classify the actions as “passing” or “clinching” actions).

First, notice that breaking information sets only shrinks the set of possible outcomes, which preserves the min/max obvious dominance inequality, and so every OSP game Γ is equivalent to a perfect information OSP game Γ' in which Nature moves once, as the first mover.¹⁹ Second, if there are two passing actions a and a' at some history h , then there are (by definition) at least two payoffs that are possible for i following each. We show that obvious dominance then implies that i must have some continuation strategy that can guarantee his top possible payoff in the continuation game following at least one of a or a' , and we can construct an equivalent game in which we replace this action with an equivalent clinching action that allows i to clinch this payoff already at h by making all such “future choices” today. This procedure can be repeated until there is at most one passing action remaining. Finally, if there remains some h such that agent i cannot clinch her favorite possible payoff at h , the game must promise i that she will never be strictly worse off by passing, which is condition 3.

Remark 1. Our proofs establish a claim stronger than the equivalence of OSP mechanisms and millipedes with greedy strategies: every obviously strategy-proof game can be transformed into a millipede game by four transformations: (i) breaking all information sets to create a perfect information game; (ii) having Nature make all of its choices at the beginning of the game; (iii) replacing continuation strategies that guarantee a payoff for an agent into a single clinching action; and (iv) Li’s pruning, in which the actions no type chooses are removed from the game tree.²⁰ These transformations are well defined even if we do not impose richness nor any other assumption on preferences. We need the richness assumption to conclude that the end-result of these transformations satisfies the third condition of the millipede definition.

Theorem 1 characterizes the entire class of obviously strategy-proof games in no-transfer environments. We have already mentioned some familiar dictatorship-like games that fit into this class (e.g., Random Priority, also known as Random Serial Dictatorship). Another example of a millipede game is sketched in Figure 2. Here, there are 100 agents $\{i, j, k_1, \dots, k_{98}\}$ and 100 objects $\{o_1, o_2, \dots, o_{100}\}$ to be assigned. The game begins with agent i being offered the opportunity to clinch o_2 , or pass to j . Agent j can then either clinch o_{99} , in which case the next mover is k_2 , or pass back to i , and so on. Now, consider the type of agent i that prefers the objects in the order of their index: $o_1 \succ_i o_2 \succ_i \dots \succ_i o_{100}$. At the very first move of the game, i is offered her second-favorite object, o_2 , even though her top choice, o_1 , is still

¹⁹That every OSP game is equivalent to an OSP game with perfect information was first pointed out in a footnote by Ashlagi and Gonczarowski (2018), which also notes that de-randomizing an OSP game leads to an OSP game. For completeness, the appendix contains the (straightforward) proofs of these statements.

²⁰The transformations (i) and (ii) are elucidated in Lemma A1, transformation (iii) in Lemma A3, and transformation (iv) in Appendix A.2.




Figure 2: An example of a millipede game.

available. The obviously dominant strategy here requires i to pass. However, if she passes, she may not be offered the opportunity to clinch her top object(s) for hundreds of moves. Further, when considering all of the possible moves of the other agents, if i passes, the game has the potential to go off into thousands of different directions, and in many of them, she will never be able to clinch better than o_2 . Thus, while passing is formally obviously dominant, fully comprehending this still requires the ability to reason far into the future of the game and perform lengthy backwards induction.

4 Simple Dominance

The upshot of the previous section is that some OSP mechanisms, such as single-unit ascending auctions and Random Priority, are indeed quite simple to play; however, the full class of millipede games is much larger, and contains OSP mechanisms that may be quite complex to play. As our analysis illustrates, the reason is that OSP relaxes the assumption that agents fully comprehend how the choices of other agents will translate into outcomes, but it still presumes that they understand how their own future actions affect outcomes. Formally, when checking obvious dominance, the min and the max are taken only over opponents' strategies, S_{-i} , fixing the agent's own strategy, S_i . Thus, while OSP guarantees that when taking an action, agents do not have to reason carefully about what their *opponents* will do, it still may require that they search deep into the game with regard to their future self, and assumes they know all of *their own* actions they will take in the future, and understand exactly how these actions will affect the set of possible outcomes. (To return to the illuminating example of chess, it presumes that at the start of the game, White knows exactly what she needs to do at any possible future configuration of the board in order to ensure a victory.)

We propose a class of simplicity standards that relax the assumption that players can analyze and plan their own actions arbitrarily far into the future of the play. The proposed conceptualization offers a way to relax the foresight assumptions embedded not only in obvious dominance but also in other game theoretic concepts. Otherwise, we maintain the approach pioneered by Li (2017b), in particular the assumptions that players cannot fully analyze the actions of other players but understand the set of possible outcomes following their own actions. In order to analyze agents who can plan for only part of the game, we need to allow the agent’s perception of the strategic situation, and hence, the planned actions—referred to as a *strategic plan* below, to distinguish from the standard game-theoretic notion of a *strategy*—to vary as the game progresses. We first formalize this idea for games of perfect information; the generalization to imperfect information will easily follow.

For each player i and node $h^* \in \mathcal{H}_i$ at which i moves, there is a set of nodes $\mathcal{H}_{i,h^*} \subseteq \{h \in \mathcal{H}_i | h \supseteq h^*\}$ that are perceived as simple from the perspective of node h^* .²¹ A **strategic plan** for agent i at node h^* is a mapping S_{i,h^*} from \mathcal{H}_{i,h^*} to actions at these nodes.²² Note that a strategic plan does not give an instruction for all continuation nodes at which i may be called to move, but rather only for those nodes in the set \mathcal{H}_{i,h^*} . Strategic plan S_{i,h^*} is **simply dominant** at node h^* if the worst possible payoff for i in the continuation game in which i follows $S_{i,h^*}(h)$ at all $h \in \mathcal{H}_{i,h^*}$ is weakly preferred by i to the best possible payoff for i in the continuation game in which i plays some other $a' \neq S_{i,h^*}(h^*)$ at h^* . A set of strategic plans, $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$, one for each node $h^* \in \mathcal{H}_i$ at which i moves, is called a **strategic collection**. A strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is **simply dominant** if all its strategic plans are simply dominant.

Given a profile of strategic collections, $S_{\mathcal{N},\mathcal{H}} = ((S_{i,h^*})_{h^* \in \mathcal{H}_i})_{i \in \mathcal{N}}$, we define a **mechanism** analogously as above, as a pair $(\Gamma, S_{\mathcal{N},\mathcal{H}})$. For any strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$, we define the **induced (global) strategy** by $S_i^*(h) = S_{i,h}(h)$, i.e., agent i ’s induced strategy at h is the action called for by the strategic plan upon reaching history h . Given a type profile $\succ_{\mathcal{N}} = (\succ_i)_{i \in \mathcal{N}}$ and a profile of strategic collections $S_{\mathcal{N},\mathcal{H}}$, we use $S_{\mathcal{N}}^*(\succ_{\mathcal{N}})$ to denote the profile of induced strategies. Note that for any $S_{\mathcal{N},\mathcal{H}}$, we can find the terminal history/outcome that is reached when the game is played according to the profile of strategic collections $S_{\mathcal{N},\mathcal{H}}$ by equivalently following the profile of induced strategies $S_{\mathcal{N}}^*(\succ_{\mathcal{N}})$. This allows us to define equivalence of mechanisms just as before, using the induced strategies: two mechanisms $(\Gamma, S_{\mathcal{N},\mathcal{H}})$ and $(\hat{\Gamma}, \hat{S}_{\mathcal{N},\mathcal{H}})$ are **equivalent** if, for every profile of types $\succ_{\mathcal{N}}$, the distribution over outcomes from the induced strategies $S_{\mathcal{N}}^*(\succ_{\mathcal{N}})$ in Γ is the same as from the induced

²¹The assumption that $\mathcal{H}_{i,h^*} \subseteq \mathcal{H}_i$ is made for simplicity; in its absence we need to endow players with beliefs of what other players will do.

²²We focus on pure strategies; the extension to mixed strategies is straightforward.

strategies $\hat{S}_{\mathcal{N}}^*(\succ_{\mathcal{N}})$ in $\hat{\Gamma}$.

This approach to simplicity takes the collection $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$ of **simple-node sets** to be a parameter of the definition. The smaller (in an inclusion sense) the set of simple nodes, the stronger is the resulting simplicity requirement. A natural requirement on the collection of simple node sets is that if an agent classifies a node $h \supset h_1$ as simple from the perspective of node h_1 then the agent continues to classify the node h as simple from the perspective of all nodes $h_2 \supseteq h_1$ such that $h \supset h_2$; we do not impose this requirement but it is satisfied in our main examples to which we turn now.

The generality of our framework allows us to embed important special cases by just varying the simple-node sets $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$. When, for any $h^* \in \mathcal{H}_i$, the set \mathcal{H}_{i,h^*} is the set of all continuation nodes of h^* at which i moves (that is, i can plan all of his future moves), we refer to any resulting simply dominant strategic collection as an **obviously dominant strategic collection**. When $\mathcal{H}_{i,h^*} = \{h^*\}$ (that is, i cannot plan any future moves), any resulting simply dominant strategic collection is called a **strongly obviously dominant strategic collection**. It is easy to see that obviously dominant strategic collections induce strategies that are obviously dominant in the sense of Li (2017b), while strongly obviously dominant strategic collections induce strategies that are strongly obviously dominant in the sense we defined in the 2016 draft of our paper.²³ Furthermore, any obviously dominant strategy naturally induces an obviously dominant strategic collection, and the same for strongly obviously dominant strategies. These are only two special cases of the concept. (Proposition 2 below shows that these special cases are in fact the extrema of the general class.) Another natural instance is when $\mathcal{H}_{i,h^*} = \{h \in \mathcal{H}_i(h^*) \mid h^* \subsetneq h' \subsetneq h \Rightarrow h' \notin \mathcal{H}_i\}$ —that is, when i can plan one move ahead but not more. We refer to simply dominant strategic collections for this set of simple nodes as **one-step foresight (OSF) dominant strategic collections**.

A strategic collection is **consistent** if $S_{i,h^*}(h) = S_{i,h}(h)$ for all $h \in \mathcal{H}_{i,h^*}$ and all $h^* \in \mathcal{H}_i$. Obviously dominant strategic collections and strongly obviously dominant strategic collections are consistent, while one-step foresight strategic collections need not be consistent.

The extension to imperfect information games is straightforward: we replace nodes h^* and h by information sets I^* and I , and replace the relationship $h \supseteq h^*$ by the relationship

²³In Section 5.2, we recall the definition of strong obvious dominance from the 2016-2018 drafts of our paper. The connection of our new approach to strong obvious dominance and to Li's obvious dominance obtains because in the definitions of these earlier dominance concepts we can replace the comparison to a strategy S'_i at the earliest point of departure with a comparison to actions other than $S_i(h)$. For instance, obviously dominant strategies can be equivalently defined as follows: a strategy S_i is *obviously dominant* if at any information set I_i the worst possible payoff to the agent from playing $S_i(I_i)$ is at least as good as the best possible payoff from any other action at I_i , where the best/worst case outcomes are determined over all possible strategies of other agents S_{-i} and all possible choices of Nature.

that I is a continuation information set of I^* (that is, I a possible information set following I^*). The key parameter of the simplicity definition is then the collection $(I_{i,I^*})_{I^* \in \mathcal{I}_i}$ of simple information sets, and the simply dominant strategic collections $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ are defined analogously to the discussion above. In the sequel, we focus on perfect information game because of the following.

Proposition 1. *Let Γ be a game of imperfect information, and consider a set of simple information sets $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$ and a corresponding simply dominant strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$. In the perfect information game in which all information sets contain exactly one history from game Γ , the induced strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is simply dominant, where each S_{i,h^*} is defined as $S_{i,h^*}(h) = S_{i,I^*}(I)$, where I is a continuation information set of I^* , $h^* \in I^*$ and $h \in I$.*

This result obtains from the analogous observation about obvious dominance—first mentioned by Ashlagi and Gonczarowski (2018) and formalized in our Lemma A1—and the first part of the following theorem.

Proposition 2. *Fix a set of simple information sets $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$. If a strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ is simply dominant for $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$, then the induced strategy $S_i^*(I^*) = S_{i,I^*}(I^*)$ is obviously dominant. Furthermore, if the induced strategy $S_i^*(I^*) = S_{i,I^*}(I^*)$ is strongly obviously dominant, then the strategic collection is simply dominant for $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$.*

This result establishes obvious dominance and strong obvious dominance as the extreme points of the class of simple dominance standards we study. It holds true because the larger (in an inclusion sense) are the sets of simple information sets the more demanding is the simple dominance requirement.

4.1 Behavioral Microfoundations

We may think of simple strategic collections as providing the right guidance to the player even if the player is confused about the action sets at non-simple nodes. We formalize this idea as follows. For expositional simplicity, we restrict attention to perfect information games. For any game Γ and collection of permutations $\tilde{\eta} = \{\eta_h\}_{h \in \mathcal{H}}$ of actions at nodes $h \in \mathcal{H}$, we construct the relabeled game $\tilde{\eta}(\Gamma)$ by permuting actions at each node h by permutation η_h ; otherwise game $\tilde{\eta}(\Gamma)$ has the same game tree as Γ and the same payoffs at terminal nodes. For instance, if $(h^*, a_1, a_2, \dots, a_k)$ is a terminal history in game Γ then it is a terminal history in game $\tilde{\eta}(\Gamma)$ and all players payoffs are the same in both games. For a set of simple nodes \mathcal{H}_{i,h^*} , we say that two games Γ and Γ' are **indistinguishable from the perspective of**

agent i at node h^* if there is a collection of permutations $\tilde{\eta} = \{\eta_h\}_{h \in \mathcal{H}}$ such that (i) η_h is an identity for all $h \in \mathcal{H}_{i,h^*}$ and (ii) $\Gamma' = \tilde{\eta}(\Gamma)$.

This preparation allows us to relate simply dominant strategic collections to the standard notion of weak dominance. We say that a strategy S_i of player i weakly dominates strategy S'_i in the continuation game beginning at h^* if following strategy S_i leads to weakly better outcomes for i than following strategy S'_i , irrespective of the strategies followed by other players. Note that here, S_i and S'_i denote full strategies in the standard game-theoretic sense of a complete contingent plan of action.

Theorem 2. *For each game Γ , agent i , preference ranking \succ_i , and collection of simple nodes $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$, the strategic plan S_{i,h^*} is simply dominant from the perspective of $h^* \in \mathcal{H}_i$ in Γ if and only if, in every game Γ' that is indistinguishable from Γ from the perspective of i at node h^* , in the continuation game starting at h^* , every strategy S_i such that $S_i(h) = S_{i,h^*}(h)$ for all $h \in \mathcal{H}_{i,h^*}$ weakly dominates any strategy S'_i such that $S'_i(h^*) \neq S_{i,h^*}(h^*)$.*

The straightforward proof is in the appendix, and—similarly to the proofs of the previous two theorems—it does not rely on our domain richness assumptions.

This theorem tells us the strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is simply dominant in Γ if and only if for every $h^* \in \mathcal{H}_i$ in every game Γ' that is indistinguishable from the perspective of node h^* every strategy S_i such that $S_i(h) = S_{i,h^*}(h)$ for all $h \in \mathcal{H}_{i,h^*}$ is weakly dominant in the continuation game starting at h^* . When the strategic collection is consistent, we can express this result equivalently in terms of simplicity of the induced global strategies $S_i(h) = S_{i,h}(h)$. When expressed in this way, this result corresponds to Li’s (2017b) microfoundation for obvious strategy-proofness.²⁴

Remark 2. Theorem 2 explains the desirability of simple dominance in terms of agents misunderstanding the game they are playing. Let us note that we could also interpret it in terms of the mechanism designer being uncertain what the agents think the game they are playing is. By designing a mechanism with simply dominant strategic collections, the designer ensures it is played as expected.

²⁴While the two results capture the same phenomenon, there is a difference between them even when restricted to OSP: Li’s (2017b) microfoundation considers a larger set of games a player might be confused between, thus—for OSP—one of his implications is formally stronger, while the other formally weaker than ours. A full analogue of Li’s result would call for a more complex formulation but it is also true in our setting, with the proof following the same steps as that of Theorem 2. Theorem 2 subsumes the microfoundation for strong obvious strategy-proofness from the 2016-2018 drafts of our paper.

5 Characterizing Simple Mechanisms

We now use our new simplicity standards to characterize what mechanisms are simple in environments both with and without transfers. For the weakest simplicity standard in the class, obvious dominance, Li (2017b) provides a characterization in environments with transfers, while we did so for no-transfer environments in Section 3 above. Thus, in this section, we focus on OSF dominance and strong obvious dominance.

5.1 One-Step Foresight Dominance

OSF dominance presumes that agents cannot plan arbitrarily far into the future of the game, but can only plan one move ahead at a time. This will eliminate the complex, yet still OSP, millipede games we saw in Section 3 in environments without transfers, while still allowing for such intuitively simple games as ascending auctions in environments with transfers. Recall that OSF dominant strategic collections obtain from simply dominant strategic collections when the simple-node sets $(\mathcal{H}_{i,h^*})_{h \in \mathcal{H}_i}$ are such that $\mathcal{H}_{i,h^*} = \{h \in \mathcal{H}_i(h^*) \mid h^* \subsetneq h' \subsetneq h \Rightarrow h' \notin \mathcal{H}_i\}$; for shorthand, we refer to any $h \in \mathcal{H}_{i,h^*} \setminus \{h^*\}$ as a **next-history** (or, next-node) at h^* . A strategic plan S_{i,h^*} is then OSF-dominant if it is simply dominant when any next-history is viewed as simple from h^* for agent i , but no other histories are simple from the perspective of h^* .

Binary allocation with transfers

One of the main applications of obvious dominance analyzed by Li (2017b) is to binary allocation with transfers. The simplest example of this is an auction of a single good, and Li (2017b) shows that in this setting, the canonical ascending (clock) auction is obviously strategy-proof, while the normal-form equivalent second-price sealed-bid auction is not. In fact, ascending auctions are not only OSP, they also satisfy the stronger property of one-step foresight simplicity. This can be seen easily by noting that the following collection of strategic plans is OSF-dominant: for any information set I_i^* such that the current price p is weakly lower than the bidder i 's value v_i : i stays In, with a plan to drop Out at any next-information set $I_i \supset I_i^*$. For any information set I_i^* such that the current price is $p > v_i$: i drops Out immediately.²⁵

Li (2017b) goes beyond just ascending clock auctions, and shows that in binary allocation settings with transfers, the class of OSP games is characterized by the class of *personal clock*

²⁵Another OSF-dominant strategic collection is: For any I_i^* such that the current price p is strictly lower than v_i : Stay In, with plan to drop Out at any next information set. For any I_i^* such that $p \geq v_i$: Drop out immediately.

auctions. As we show next, personal clock auctions are also OSF-simple, and so, surprisingly, any OSP-implementable social choice rule is also implementable in OSF-dominant strategic collections.

Formally, we follow Li (2017b) and define binary allocation problems with transfers as follows: $\mathcal{Y} \subseteq 2^{\mathcal{N}}$ is a set of possible allocations and $w \equiv (w_i)_{i \in \mathcal{N}} \in \mathcal{W}^{\mathcal{N}}$ is a profile of transfers. The set of outcomes is thus $\mathcal{X} = \mathcal{Y} \times \mathcal{W}^{\mathcal{N}}$. To be consistent with Li (2017b), in this section, we denote agents types by $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$, where $0 \leq \underline{\theta}_i < \bar{\theta}_i$, and assume that agents have quasilinear preferences, where $u_i(\theta_i, y, w) = 1_{i \in y} \theta_i + w_i$ denotes agent i 's utility from outcome (y, w) when she has type θ_i .²⁶ For a (perfect-information) game Γ , define outcome functions g such that $g_y(\bar{h}) \in \mathcal{Y}$ is the allocation at terminal history \bar{h} , and $g_{w,i}(\bar{h}) \in \mathcal{W}$ is the transfer to agent i at \bar{h} . The following definition of a personal clock auction is adapted from Li (2017b). Note that the game is deterministic, i.e., there are no moves by Nature.²⁷

Γ is a **personal clock auction** if, for every $i \in \mathcal{N}$, at every earliest history at which i moves h_i^* , either:

In Transfer Falls: there exists a fixed transfer $\bar{w}_i \in \mathcal{W}$, a going transfer $\tilde{w}_i : \{h_i : h_i^* \subseteq h_i\} \rightarrow \mathcal{W}$ and a set of “quitting actions” A^q such that

- For all terminal $\bar{h} \supset h_i^*$, either (i) $i \notin g_y(\bar{h})$ and $g_{w,i}(\bar{h}) = \bar{w}_i$ or (ii) $i \in g_y(\bar{h})$ and $g_{w,i}(\bar{h}) = \min\{\tilde{w}_i(h_i) : h_i^* \subseteq h_i \subsetneq \bar{h}\}$.
- If $\bar{h} \supsetneq (h, a)$ for some $h \in \mathcal{H}_i$ and $a \in A^q$, then $i \notin g_y(\bar{h})$.
- $A^q \cap A(h_i^*) \neq \emptyset$
- For all $h'_i, h''_i \in \{h_i : h_i^* \subseteq h_i\}$:
 - If $h'_i \subsetneq h''_i$, then $\tilde{w}_i(h'_i) \geq \tilde{w}_i(h''_i)$
 - If $h'_i \subsetneq h''_i$, $\tilde{w}_i(h'_i) > \tilde{w}_i(h''_i)$ and there is no h'''_i such that $h'_i \subsetneq h'''_i \subsetneq h''_i$, then $A^q \cap A(h''_i) \neq \emptyset$
 - If $h'_i \subsetneq h''_i$ and $\tilde{w}_i(h'_i) > \tilde{w}_i(h''_i)$, then $|A(h'_i) \setminus A^q| = 1$
 - If $|A(h'_i) \setminus A^q| > 1$, then there exists $a \in A(h'_i)$ such that, for all $\bar{h} \supseteq (h'_i, a)$, $i \in g_y(\bar{h})$.

²⁶Note that Li (2017b) allows for a continuum of transfers and types. Our simplicity standards extend straightforwardly to this environment.

²⁷In light of our notion of equivalent mechanisms and Proposition 1, the presentation below is a simplification of Definition 15 of Li (2017b): for any personal clock auction that satisfies Definition 15 of Li (2017b), there is an equivalent mechanism that satisfies the definition below. (This statement is valid for original Li’s (2017b) definition and for the one presented in the corrigendum available from his website).

or, **Out Transfer Falls:**

- As above, but replace every instance of “ $i \in g_y(\bar{h})$ ” with “ $i \notin g_y(\bar{h})$ ” and vice-versa.

Intuitively, a personal clock auction is a generalization of classic ascending/descending auctions. In a standard ascending auction for a single good, there is a single price for all agents; at each history, agents choose from one of two actions, either “quit” or “continue”; when an agent quits, she does not win the object (i.e., is out of the allocation), and receives a transfer of zero. As Li (2017b) discusses, personal clock auctions generalize this procedure in several ways: agents may face different prices (“clocks”); at any history, there may be multiple quitting or multiple continuing actions; when an agent quits, her transfer need not be zero; some agents may face In-Transfer Falls while others face Out-Transfer Falls (a two-sided clock auction). The key restrictions are that the clock for each agent can only go in one direction (i.e., either In-Transfer Falls or Out-Transfer Falls), and, whenever the transfer an agent faces strictly changes, she must be offered an opportunity quit. But, these restrictions also ensure that there is an OSF-dominant strategic plan at any h_i . In particular, we have the following result.

Theorem 3. *In binary allocation settings with transfers, every OSF-simple mechanism is equivalent to a personal clock auction. Furthermore, every personal clock auction is OSF-simple.*

Because our Proposition 2 implies that any OSF-simple mechanism is also OSP, the first part of the theorem readily follows from Li’s (2017b) result that the class of OSP mechanisms is characterized by personal clock auctions. That personal clock auctions are OSF-simple is straightforward to check, establishing the second part of the theorem.

Environments without transfers

In Section 3, we saw that in no-transfer environments, some millipede games, while obviously strategy-proof, could still be quite complex (e.g., Figure 2). Imposing the stronger standard of OSF-simplicity eliminates such complex millipede games, and leaves only games that are monotonic in the following sense: a game Γ is **monotonic** if, for any agent i and any histories $h \subsetneq h'$ such that $i_h = i$, $i_{h'} = i$ or h' is terminal, and $i_{h''} \neq i$ for any h'' such that $h \subsetneq h'' \subsetneq h'$, either (i) $C_i(h) \subseteq C_i(h')$ or (ii) $P_i(h) \setminus C_i(h) \subseteq C_i(h')$. In words, this says that at any next-history for i , she is offered to clinch either (i) everything she could have clinched at her previous move or (ii) everything that was possible, but not clinchable at her previous move.

Theorem 4. *In environments without transfers, every OSF-simple game is equivalent to a monotonic millipede game. Furthermore, every monotonic millipede game is OSF-simple.*

From the perspective of an agent playing in a game, monotonic games seem particularly simple: each time an agent is called to move, she knows that if she chooses to pass (i.e., not clinch), at her next move, she will either be able to clinch everything she is offered to clinch currently, or she will be able to clinch her top remaining choice. On the other hand, in a non-monotonic game such as that in Figure 2, an agent’s possible clinching options at future moves may be strictly worse for her for many moves in the future, before eventually being re-offered what she was able to clinch in the past. If agents are unable to plan far ahead in the game tree, it may be difficult to recognize that passing is obviously dominant; in a monotonic game, however, agents only need to be able to plan at most one step at a time to recognize that passing is a dominant choice.

Remark 3. The idea of monotonicity extends to our general environment but we need to replace inclusions by dominance relations between sets. We say that, for an agent i , set A of payoffs is **dominated** by set B of payoffs, and write $A \leq_i B$, if for every payoff $a \in A$ there is a payoff $b \in B$ such that $b \succeq a$. We then say that a game Γ is **monotonic** if, for any agent i and any histories $h \subsetneq h'$ such that $i_h = i$, $i_{h'} = i$ or h' is terminal, and $i_{h''} \neq i$ for any h'' such that $h \subsetneq h'' \subsetneq h'$, either (i) $C_i(h) \leq_i C_i(h')$ or (ii) $P_i(h) \setminus C_i(h) \leq_i C_i(h')$. Personal clock auctions from the previous subsection are monotonic in this sense, while in settings in which no payoff is dominated by another—that is, in settings which we referred to as settings without transfers—these dominance relations reduce to inclusions of the sets of payoffs and we recover the monotonicity definition used in Theorem 4.

5.2 Strong Obvious Dominance and Price Mechanisms

In light of Theorem 2, the strongest simplicity standard in our class is strong obvious dominance. To remind, a strategy S_i is **strongly obviously dominant** if, for any other strategy S'_i , starting at any earliest point of departure h between S_i and S'_i , the worst possible outcome from following S_i is weakly better than the best possible outcome following S'_i , where the best and worst cases are taken over all future actions of other agents (including Nature) and all future actions of agent i . In the framework of strategic plans/collections, strong obvious dominance obtains from simply dominant strategic collections when the set of simple nodes from the perspective of h^* is $\mathcal{H}_{i,h^*} = \{h^*\}$. If a game Γ admits a profile of strongly obviously dominant strategies, we say that it is **strongly obviously strategy-proof (SOSP)**. Random Priority is SOSP, but the millipede game depicted in Figure 2 is not. Thus, SOSP mechanisms further delineate the class of games that are simple to play, by eliminating the more complex millipede games that may require significant forward-looking behavior and backward induction.

Strong obvious strategy-proofness has several appealing features that capture the idea of a game being simple to play. Since SOSP strengthens OSP by looking at the worst/best case outcomes for i over all possible future actions that could be taken by i 's opponents and agent i herself, a strongly obviously dominant strategy is one that is weakly better than all alternative strategies even if the agent is concerned that she might tremble in the future or has time-inconsistent preferences. Further, SOSP games can be implemented so that each agent is called to move at most once. We can actually show a stronger result that highlights the simplicity of SOSP games: in any SOSP game, each agent can have at most one history at which her choice of action is payoff-relevant. Formally, we say a history h at which agent i moves is payoff-irrelevant for this agent if i receives the same payoff at all terminal histories $\bar{h} \supset h$; if i moves at h and this history is not payoff-irrelevant, then it is **payoff-relevant** for i . The definition of SOSP and richness of the preference domain give us the following.

Lemma 1. *Along each path of an SOSP game that is on the path of the greedy strategies for some type profile, there is at most one payoff-relevant history for each agent.*

This result allows us to further conclude that, for a given game path, the unique payoff-relevant history (if it exists) is the first history at which an agent is called to move.²⁸ While an agent might be called to act later in the game, and her choice might influence the continuation game and the payoffs for other agents, it cannot affect her own payoff.

Building on Lemma 1, we show that SOSP effectively implies that agents—in a sequence—are faced with choices from personalized menus (e.g., in allocation with transfers this may be menus of object-price pairs). At the typical payoff-relevant history an agent is offered a menu of payoffs that she can clinch, she selects one of the alternatives from the menu, and she is never called to move again. More formally, we say that Γ is a **sequential price game** if it is a perfect-information game in which Nature moves first (if at all). The agents then move sequentially, with each agent called to play at most once. The ordering of the agents and the sets of possible outcomes at each history are determined by Nature's action and the actions taken by earlier agents. As long as there are either at least three distinct undominated payoffs possible for the agent who is called to move or there is exactly one such payoff, the agent can clinch any of the possible payoffs, while at the same time also selecting a message from a pre-determined set of messages. When exactly two undominated payoffs are possible for the agent who moves, the agent can be faced with either (i) a choice between them (clinching and selecting an accompanying message), or, (ii) a choice between clinching one of these payoffs (and selecting an accompanying message) and passing (with

²⁸The on-path restriction is not needed if we consider the class of “pruned” games in the sense of Li (2017b); cf. Appendix A.2 for how pruning works with OSP, SOSP, and other simplicity concepts.

no message); note that (ii) does not allow the agent to clinch the other payoff. Also note that, similarly to OSP games, in a sequential price game, an agent may have several actions that clinch the same payoff for herself; the purpose of including messages (in addition to selecting a payoff) is to capture this feature. In environments without transfers, sequential price games resemble sequential dictatorships.

Theorem 5. *Every strongly obviously strategy-proof mechanism (Γ, S_N) is equivalent to a sequential price mechanism with the greedy strategy. Every sequential price mechanism with the greedy strategy is strongly obviously strategy-proof.*

Theorem 5 applies to a wide array of environments. For instance, in an object allocation model without transfers, every SOSP mechanism resembles a sequential dictatorship in which agents are called sequentially and offered some subsets of objects that they can clinch; they pick their most preferred object and leave the game (the difference between a sequential price game and a sequential dictatorship is that at an agent’s turn, she need not be offered all still-available objects). In a binary allocation setting with a single good and transfers, each agent is approached one at a time, and given a take-it-or-leave-it (TIOLI) offer of a price at which she can purchase the good; if an agent refuses, the next agent is approached, and given a (possibly different) TIOLI offer, etc. If there are multiple objects for sale, each agent is offered a menu consisting of several bundles of objects with associated transfers, and selects her most preferred option from the menu. These are only a few examples covered by Theorem 5; the result holds for any environment that satisfies the richness assumption from Section 2.

6 Random Priority

As an application of our study of simplicity, we show that OSP can be combined with natural fairness and efficiency axioms to provide a characterization of the popular Random Priority (RP) mechanism. In Random Priority, first Nature selects an ordering of agents, and then each agent moves in turn and chooses her favorite object among those that remain available given previous agents’ choices. Random Priority succeeds on three important design dimensions: it is simple to play, efficient, and fair.²⁹ However, this is only a partial explanation of its success, as to now, it has remained unknown whether there exist other such mechanisms,

²⁹Pareto efficiency and fairness of RP have been recognized at least since Abdulkadiroğlu and Sönmez (1998) (see Bogomolnaia and Moulin (2001) for analysis of more demanding efficiency concepts), while Li (2017b) established OSP of RP. It is easy to see that the standard extensive-form implementation of RP also satisfies all of our more demanding simplicity requirements.

and, if so, what explains the relative popularity of RP over these alternatives.³⁰ Theorem 6 provides an answer to this question: not only does Random Priority have good efficiency, fairness, and incentive properties, it is the *only* mechanism that does so, thus explaining the widespread popularity of Random Priority in practice.

We consider a canonical object allocation model with single-unit demand, a special case of our general framework. There is a set \mathcal{N} of agents, a set of objects, also of cardinality $|\mathcal{N}|$, and global outcomes are bijections between agents and objects. Each agent has a strict preference ranking \succ_i over the objects. Our efficiency concept is Pareto efficiency: an outcome is **Pareto efficient** when no other outcome is weakly preferred by all participants and strictly preferred by at least one; a mechanism $(\Gamma, S_{\mathcal{N}})$ is **Pareto efficient** if it generates Pareto efficient outcomes for all Nature’s choices and agents’ types.³¹ Our fairness concept is symmetry: a mechanism $(\Gamma, S_{\mathcal{N}})$ is **symmetric** if, for any two agents $i, j \in \mathcal{N}$, the outcome distribution of the mechanism does not change when we transpose the preference rankings of i and j and at the same time transpose the objects the two agents obtain. Informally, the outcome of the mechanism would not change if i played the role of j and vice versa.³² The symmetry condition fails in a serial dictatorship in which player 1 chooses first among all outcomes and then player 2 chooses among all remaining outcomes: if they have the same most preferred object then 1 obtains this object in the original serial dictatorship but not in the transposed one. Random Priority orders the agents randomly, and in effect the probability agent 1 obtains the preferred object is the same before and after the transposition.

Theorem 6. *An obviously strategy-proof mechanism is symmetric and Pareto efficient if and only if it is equivalent to Random Priority.*

That RP is obviously strategy-proof was recognized by Li (2017b), and its Pareto efficiency and symmetry is known at least since Abdulkadiroğlu and Sönmez (1998). The converse is new. A key step in the proof is our construction of a bijection between permutations of any deterministic Pareto-efficient millipede and permutations of serial dictatorships such that the outcomes of the permuted millipede and permuted serial dictatorship are exactly the same. Applying a permutation of agents $\sigma : \mathcal{N} \rightarrow \mathcal{N}$ to a serial dictatorship means that we use σ to change the order in which agents make their choices; similarly, applying

³⁰Bogomolnaia and Moulin (2001) provide a characterization of RP in the special case of $|\mathcal{N}| = 3$, but their result does not extend to larger markets; Liu and Pycia (2011) provide a characterization using asymptotic versions of standard axioms in replica economies as the market size grows to infinity.

³¹Because our simplicity axiom will be obvious dominance, $S_{\mathcal{N}}$ here denotes a profile of strategies in the standard game-theoretic sense (rather than strategic plans).

³²We formalize the concept of the role in the appendix. Because any permutation can be decomposed into a composition of transpositions, we can equivalently state the symmetry property as $\sigma^{-1} \circ (\Gamma, S_{\mathcal{N}}) \circ \sigma = (\Gamma, S_{\mathcal{N}})$ for all permutations $\sigma : \mathcal{N} \rightarrow \mathcal{N}$.

the permutations σ to a millipede means that agent i is given the moves of agent $\sigma(i)$ (see appendix for more formal treatment). The bijection idea was first employed by Abdulkadiroğlu and Sönmez (1998), and has since been used by several others (e.g., Pathak and Sethuraman (2011) and Carroll (2014)). Our construction of the bijection is highly involved and very different from the bijections of the earlier literature. In the construction, we rely on the properties of the millipedes established by us, and on the properties of Pareto efficient OSP mechanisms subsequently obtained by Bade and Gonczarowski (2017). The bijection argument only proves the special case of the theorem restricted to mechanisms that take the form of a uniform randomization over permutations of a deterministic Pareto-efficient millipede, and another key step of the proof is showing that every relevant symmetric mechanism is equivalent to a lottery over such uniform randomizations. We provide details in the appendix.

7 Conclusion

We study the question of what makes a game “simple to play”, and introduce a general class of simplicity standards that vary the foresight abilities required of agents in extensive-form imperfect-information games. We consider agents that form a strategic plan only for a limited horizon in the continuation game, though they may update these plans as the game progresses and the future becomes the present. The least restrictive simplicity standard include in our class is Li’s (2017b) obvious strategy-proofness, which presumes agents have unlimited foresight of their own actions, while the strongest, strong obvious strategy-proofness, presumes no foresight. For each of these standards, as well as an intermediate standard of one-step-foresight simplicity, we provide characterizations of simple mechanisms in various environments with and without transfers, and show that our simplicity standards delineate classes of mechanisms that are commonly observed in practice. Among these results, we show that Li’s characterization of OSP mechanisms as personal clock auctions can be strengthened to OSF, and that SOSP mechanisms are equivalent to price mechanisms, which are ubiquitous in practice.³³ Finally, in the context of object allocation without transfers, we provide an explanation for the popularity of Random Priority by showing that it is the essentially unique mechanism that is OSP, efficient, and symmetric.

Our results contribute to the understanding of the fundamental trade-off between simplicity of mechanisms and the ability to implement other social objectives, such as efficiency. In environments with transfers, Vickrey (1961), Riley and Samuelson (1981), Myerson (1981),

³³Even on eBay, which began as an auction website, Einav et al. (2018) document a dramatic shift in the 2000s from auctions to posted prices as the predominant selling mechanism on the platform.

Manelli and Vincent (2010), and Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) show that the efficiency achieved with Bayesian implementation can be replicated in dominant strategies; thus the accompanying increase in simplicity may come without efficiency costs. Li (2017b) and our paper advance this insight yet further and establish that obviously strategy-proof and one-step-foresight simple mechanisms can also implement efficient outcomes. At the same time, strong obvious dominance is more restrictive, and will more severely limit the class of implementable objectives. In environments with transfers, SOSP will generally preclude efficiency,³⁴ though it is important to note that in environments without transfers, even SOSP mechanisms—serial dictatorships—can achieve efficient outcomes. Combining our results with the mechanism equivalence analysis of Pycia (2017) allows us to conclude that, in single-unit demand allocation problems without transfers, the restriction to strongly obvious strategy-proof mechanisms allows the designer to achieve virtually the same efficiency and many other objectives as those achievable in merely strategy-proof mechanisms. Thus in many environments, simplicity entails no efficiency loss. In other environments, the trade-off between simplicity and efficiency is more subtle.

Our work is complementary to the experimental literature on how mechanism participants behave and what elements of design enable them to play equilibrium strategies, cf. e.g. Kagel et al. (1987) and Li (2017b). While this literature identifies implementation features that facilitate play and confirms that obviously strategy-proof mechanisms are indeed simpler to play than merely strategy-proof mechanisms, and that strongly obviously strategy-proof mechanisms are easier still and nearly all participants play them as expected,³⁵ our general theory of simplicity opens new avenues for experimental investigations. For instance, one may define the simplicity level of a game in terms of the smallest (in an inclusion sense) set of histories that an agent must see as simple in the sense of Section 4 in order to play the equilibrium strategy correctly. One may similarly measure the sophistication of experimental subjects.

In sum, the sophistication of agents may vary across applications, and so it is important to have a range of simplicity standards.³⁶ For sophisticated agents, a weaker simplicity standard

³⁴For instance, when we want to allocate an object to the highest value agent in an environment with transfers in which there are at least two agents and agents' values are drawn iid from among at least three values, an impossibility result obtains: no SOSP and efficient mechanism exists. This is implied by our sequential price mechanism characterization (Theorem 5). On the other hand, Armstrong (1996) shows that posted prices achieve good revenues when bundling allows the seller to equalize the valuations of buyers, and Chawla, Hartline, Malec, and Sivan (2010) and Feldman, Gravin, and Lucier (2014) show that sequential price mechanisms achieve decent revenues even without the bundling/equalization assumption.

³⁵See Bo and Hakimov (2019). For tests of the first claim see also Li (2017b) and Breitmoser and Schweighofer-Kodritsch (2019).

³⁶Our theory provides a partial ordering on many simplicity standards but not all; for an example of a simplicity standard not encompassed by our theory, see Börgers and Li (2019).

ensures they play the intended strategies, allowing in principle the designer more flexibility on other objectives; however, for less sophisticated agents, a stronger standard of simplicity may need to be imposed, with the potential limitations on flexibility. Understanding the simplicity of games and the simplicity-flexibility tradeoff requires an adaptable approach to thinking about simplicity. This paper puts forth one such proposal, though there is much work still to be done in fully exploring this trade-off and testing various simplicity standards empirically.

A Proofs

A.1 Preliminary Definitions

Before proceeding with the main proofs, we first define the concepts of possible, guaranteeable, and clinchable outcomes/actions more formally. Fix a game Γ . Let $S = (S_i)_{i \in \mathcal{N}}$ denote a strategy profile for the agents. Let $\omega := (\omega(h))_{\{h \in \mathcal{H}: \text{Nature moves at } h\}}$ denote one particular realization of Nature's moves through the game, where $\omega(h) \in A(h)$ is the action taken by Nature at a history h at which Nature is called to move. Define $z(h, S, \omega) \in \mathcal{X}$ as the unique final outcome that obtains at the terminal history \bar{h} that is reached when play starts at some h and proceeds according to (S, ω) .

We first discuss the distinction between types of payoffs (possible vs. guaranteeable) and then the distinction between types of actions (clinching actions vs. passing actions). Recall that agents may be indifferent between several outcomes. For any outcome $x \in \mathcal{X}$, let $[x]_i = \{y \in \mathcal{X} : y \sim_i x\}$ denote the x -indifference class of agent i , and define

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

to be the possible indifference classes that may obtain for agent i starting at history h if she follows strategy S_i . If there exists some S_i such that $[x]_i \in X_i(h, S_i)$, then we then we say that $[x]_i$ is **possible** for i at h . If, further, there exists some S_i such that $X_i(h, S_i) = \{[x]_i\}$, then we say $[x]_i$ is **guaranteeable** for i at h . Let

$$\begin{aligned} P_i(h) &= \{[x]_i : \exists S_i \text{ s.t. } [x]_i \in X_i(h, S_i)\} \\ G_i(h) &= \{[x]_i : \exists S_i \text{ s.t. } X_i(h, S_i) = \{[x]_i\}\} \end{aligned}$$

be the sets of indifference classes that are possible and guaranteeable at h , respectively.³⁷ Note that $G_i(h) \subseteq P_i(h)$, and the set $P_i(h) \setminus G_i(h)$ is the set of indifference classes that are

³⁷Note that $P_i(h)$ and $G_i(h)$ are well-defined even if $i_h \neq i$, i.e., even if i is not the agent who moves at h .

possible at h , but are not guaranteeable at h .

Last, we define a distinction between two kinds of actions: clinching actions and passing actions. Let $i_h = i$ be the agent who is to act at a history h . Using our notational convention that (h, a) denotes the history obtained by starting at h and following action a , the set $P_i((h, a))$ is the set of payoffs that are possible for i if she takes action a at h . If $P_i((h, a)) = \{[x]_i\}$, then we say that action $a \in A(h)$ **clinches** payoff x for i . If an action a clinches x for i , we call a a **clinching action**. Note that there can be multiple actions in $A(h)$ that clinch the same payoff x for i . Any action of an agent that is not a clinching action is called a **passing action**. Let $C_i(h)$ denote the set of payoffs that are clinchable for i at h ; that is,

$$C_i(h) = \{[x]_i : \exists a \in A(h) \text{ s.t. } P_i((h, a)) = \{[x]_i\}\}.$$

Note that this definition of $C_i(h)$ presumes that agent i is called to play at history h . If \bar{h} is a terminal history, then no agent is called to play and there are no actions. However, it will be useful in what follows to define $C_i(\bar{h}) = \{[x]_i\}$ for all i , where x is the unique outcome associated with the terminal history \bar{h} .

We also remind the reader of two additional pieces of notation that were introduced in Section 3:

$$C_i^{\subseteq}(h) = \{[x]_i : [x]_i \in C_i(h') \text{ for some } h' \subseteq h \text{ s.t. } i_{h'} = i\}$$

$$C_i^{\subset}(h) = \{[x]_i : [x]_i \in C_i(h') \text{ for some } h' \subsetneq h \text{ s.t. } i_{h'} = i\}.$$

In words, $C_i^{\subseteq}(h)$ is the set of payoffs that i can clinch at some subhistory of h , and $C_i^{\subset}(h)$ is the set of payoffs that i can clinch at some *strict* subhistory of h . Note that the definition of $C_i(h)$ implicitly presumes that $i_h = i$, i.e., i moves at h ; however, $P_i(h)$, $C_i^{\subseteq}(h)$ and $C_i^{\subset}(h)$ are defined for any h , whether i is the agent who moves at h or not.

A.2 Pruning Principle

Li (2017b) introduced the following pruning principle for obvious dominance. Given a game Γ and strategy profile $(S_i(\succ_i))_{i \in \mathcal{N}}$, the **pruning** of Γ with respect to $(S_i(\succ_i))_{i \in \mathcal{N}}$ is a game Γ' that is defined by starting with Γ and deleting all histories of Γ that are never reached for any type profile. Then, the **pruning principle** says that if $(S_i(\succ_i))_{i \in \mathcal{N}}$ is obviously dominant for Γ , the restriction of $(S_i(\succ_i))_{i \in \mathcal{N}}$ to Γ' is obviously dominant for Γ' , and both games result in the same outcome. Thus, for any OSP mechanism, we can find an equivalent OSP pruned mechanism. For strong obvious dominance the pruning principle remains valid.

For all simple dominance concepts we study, a modified pruning principle is valid: we only delete histories that are never simple on the path of the game for any type profile.

A.3 Proof of Theorem 1

We break Theorem 1 into two propositions. We start by proving that millipede games are OSP (Proposition A1), and then prove that every OSP game is equivalent to a millipede game (Proposition A2).

Proposition A1. *Millipede games with greedy strategies are obviously strategy-proof.*

Proof. Let Γ be a millipede game. Recall that the *greedy strategy* for any agent i is defined as follows: for any history h at which i moves, if i can clinch her top payoff in $P_i(h)$, then $S_i(\succ_i)(h)$ instructs i to follow an action that clinches this payoff; otherwise, i passes at h .³⁸

We now show that it is obviously dominant for all agents to follow a greedy strategy. Consider some profile of greedy strategies $(S_i(\cdot))_{i \in \mathcal{N}}$. 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 $x \succeq_i y$ for all $y \in X'$ (note that we use our standard convention whereby a payoff x represents the entire indifference class to which x belongs, and so $Top(\succ_i, X')$ is a singleton). 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 . It is clear that if $Top(\succ_i, C_i(h)) = Top(\succ_i, P_i(h))$, then the greedy action of clinching the top payoff is obviously dominant at h . 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 .

Assume that there exists a history h that is on the path of play for type \succ_i when she follows the greedy strategy and $Top(\succ_i, C_i(h)) \neq Top(\succ_i, P_i(h))$, yet passing is not obviously dominant at h ; further, let h be any earliest such history for which this is true. 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 payoff from passing (and following $S_i(\succ_i)$ in the future). Since passing is not obviously dominant, it must be that $x_W(h) \not\succeq_i x_C(h)$.

First, note that $x_W(h) \succeq_i x_W(h')$ for all $h' \subsetneq h$ such that $i_{h'} = i$. Since passing is obviously dominant at all such h' , we have $x_W(h') \succeq_i x_C(h')$, and together, these imply that $x_W(h) \succeq_i x_C(h')$ for all such 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 $x' \succ_i x_C(h) \succ_i$

³⁸There may be multiple ways for i to clinch the same payoff x at h , and further, x may in principle still be possible/guaranteeable if i passes at h . Our goal is simply to prove the existence of at least one obviously dominant strategy for i .

$x_W(h)$.³⁹ The above implies that $x' \succ_i x_C(h) \succ_i x_C(h')$ for all $h' \subsetneq h$ such that $i_{h'} = i$. Let $X_0 = \{x' : x' \in P_i(h) \text{ and } x' \succ_i x_C(h)\}$. 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 $x_1 \succ_i x_2 \succ_i \dots \succ_i x_M$.

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 payoff $x_W(h)$. For every $x_m \in X_0$, let h_m denote the history on this path at which x_m becomes impossible for i .⁴⁰ Note that because i is ultimately receiving payoff $x_W(h)$, such a history h_m exists for all $x_m \in X_0$.⁴¹ Let $\hat{h} = \max\{h_1, h_2, \dots, h_M\}$ (ordered by \subset); in words, \hat{h} is the earliest history at which everything in X_0 is no longer possible. Further, let $\hat{h}_{-m} = \max\{h_1, \dots, h_{m-1}\}$, i.e., \hat{h}_{-m} is the earliest history at which all payoffs strictly preferred to x_m are no longer possible.

Claim A1. For all $x_m \in X_0$ and all $h' \subseteq \bar{h}$, we have $x_m \notin C_i(h')$.

Proof. First, note that $x_m \notin C_i(h')$ for any $h' \subseteq h$ by construction. We will show that $x_m \notin C_i(h')$ at any $\bar{h} \supseteq h' \supset h$ as well. Start by considering $m = 1$, and assume $x_1 \in C_i(h')$ for some $\bar{h} \supseteq h' \supset h$. By definition, $x_1 = \text{Top}(\succ_i, P_i(h))$; since $h' \supset h$ implies that $P_i(h') \subseteq P_i(h)$, we have that $x_1 = \text{Top}(\succ_i, P_i(h'))$ as well. Since $x_1 \in C_i(h')$ by supposition, greedy strategies direct i to clinch x_1 , which contradicts that she receives $x_W(h)$.⁴²

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

Case (i): $h' \subset \hat{h}_{-m}$. This is the case where x_m is clinchable while there is some strictly preferred payoff $x_{m'} \succ_i x_m$ that is still possible. Since $x_{m'}$ becomes impossible at \hat{h}_{-m} and is previously unclinchable, by definition of a millipede game, so $x_m \in C_i(\hat{h}_{-m})$. 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): $h' \supseteq \hat{h}_{-m}$. 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 $x_W(h)$. \square

³⁹At least one such x' exists by the assumption that $\text{Top}(\succ_i, C_i(h)) \neq \text{Top}(\succ_i, P_i(h))$, though there in general may be multiple such x' .

⁴⁰Recall from the main text that we say a payoff x becomes impossible for i at h if it was possible for all prior histories at which i moves, and is no longer possible at h .

⁴¹It is possible that h_m is a terminal history.

⁴²Recall that for terminal histories h , we define $C_i(h) = \{x\}$, where x is the unique payoff associated with the terminal history. Thus, if h' is a terminal history, then i receives payoff x_1 , which also contradicts that she receives payoff $x_W(h)$.

To finish the proof, again let $\hat{h} = \max\{h_1, h_2, \dots, h_M\}$ and let \hat{x} be a payoff that becomes impossible at \hat{h} . The claim shows that \hat{x} is not clinchable at any $h' \subseteq \hat{h}$. Therefore, by part 3 in the definition of a millipede game, $x_C(h) \in C_i(\hat{h})$. Since $x_C(h)$ is the best possible remaining payoff at \hat{h} , greedy strategies direct i to clinch $x_C(h)$, which contradicts that she receives $x_W(h)$.⁴³ \square

Proposition A2. *Every obviously strategy-proof mechanism (Γ, S_N) is equivalent to a millipede game with greedy strategies.*

Proof. Given Li’s pruning principle (see Subsection A.2), Proposition A2 follows directly from Lemmas A1, A3, and A4 that we state and prove below.⁴⁴

Lemma A1. *Every OSP game is equivalent to an OSP game with perfect information in which Nature moves at most once, as the first mover.*

Proof. Ashlagi and Gonczarowski (2018) briefly mention this result in a footnote; here, we provide the straightforward proof for completeness. We first show that every OSP game is equivalent to an OSP game with perfect information. 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 . Thus, the worst-case outcome from any action (weakly) increases in Γ' , while the best-case outcome (weakly) decreases. Thus, if there is an obviously dominant strategy in Γ , following the analogous strategy in Γ' continues to be obviously dominant. Hence, Γ' is obviously strategy-proof and equivalent to Γ .

We now show that every OSP game is equivalent to a perfect-information OSP game in which Nature moves once, as the first mover. Consider a game Γ , which, by the previous paragraph, we can assume has perfect information. Let $\mathcal{H}_{\text{nature}}$ be the set of histories h at which Nature moves in Γ . Consider a modified game Γ' in which at the empty history Nature chooses actions from $\times_{h \in \mathcal{H}_{\text{nature}}} A(h)$. After each of Nature’s initial moves, we replicate the original game, except at each history h at which Nature is called to play, we delete Nature’s move and continue with the subgame corresponding to the action Nature chose from $A(h)$

⁴³If \hat{h} is a terminal history, then we make an argument analogous to footnote 42 to reach the same contradiction.

⁴⁴We actually prove a slightly stronger statement, which is that every OSP game is equivalent to a millipede game that satisfies the following additional property: for all i , all h at which i moves, and all $x \in G_i(h)$, there exists an action $a_x \in A(h)$ that clinches x (see Lemma A3 below).

at \emptyset . Again, note that for any agent i and history h at which i is called to act, the support of possible outcomes at 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). Using reasoning similar to the previous paragraph, we conclude that Γ' is obviously strategy-proof, and Γ and Γ' are equivalent. \square

Lemma A2. *Let Γ be an obviously strategy-proof game of perfect information that is pruned with respect to the obviously dominant strategy profile $(S_i(\succ_i))_{i \in \mathcal{N}}$. Consider a history h where agent $i_h = i$ is called to move. There is at most one action $a^* \in A(h)$ such that $P_i((h, a^*)) \not\subseteq G_i(h)$.*

Proof. For any history h , let $PnG_i(h) = P_i(h) \setminus G_i(h)$ (where “**PnG**” is shorthand for “possible but not guaranteeable”). Now, consider any h at which i moves, and assume that at h , there are (at least) two such actions $a_1^*, a_2^* \in A(h)$ as in the statement. We first claim that $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \emptyset$, where $h_1^* = (h, a_1^*)$ and $h_2^* = (h, a_2^*)$. Indeed, if not, then let x be a payoff in this set. By pruning, some type \succ_i is following some strategy such that $S_i(\succ_i)(h) = a_1^*$ that results in a payoff of x at some terminal history $\bar{h} \supset (h, a_1^*)$. Note that $Top(\succ_i, P_i(h)) \neq x$, because otherwise a_1^* would not be obviously dominant for this type (since $x \notin G_i(h)$ and $x \in P_i(h_2^*)$). Thus, let $Top(\succ_i, P_i(h)) = y$. Note that $y \notin G_i(h)$ (or else it would not be obviously dominant for type \succ_i to play a strategy such that x is a possible payoff), and it is without loss of generality to assume that $Top(\succ_i, P_i(h) \setminus \{y\}) = x$.⁴⁵ Further, note that $y \in P_i(h_1^*)$ and $y \notin P_i(h_2^*)$. To see the former, note that if $y \notin P_i(h_1^*)$, then a_1^* is not obviously dominant for type \succ_i , which contradicts that $S_i(\succ_i)(h) = a_1^*$; given the former, if $y \in P_i(h_2^*)$, then once again a_1^* would not be obviously dominant for type \succ_i . Now, again by pruning, there must be some type \succ'_i such that $S_i(\succ'_i)(h) = a_2^*$ that results in payoff x at some terminal history $\bar{h} \supset (h, a_2^*)$. By similar reasoning as previously, $Top(\succ'_i, P_i(h)) \neq x$, and so $Top(\succ'_i, P_i(h)) = z$ for some $z \in P_i(h_2^*)$. Since $y \notin P_i(h_2^*)$, we have $z \neq y$, and we can as above conclude that $z \notin G_i(h)$. Similarly to footnote 45, it is without loss of generality to consider the type \succ''_i that ranks y immediately after z . Note that, for this type, no action $a \neq a_2^*$ can obviously dominate a_2^* (since $z \notin G_i(h)$). Further, a_2^* itself is not obviously dominant for this type, since the worst case from a_2^* is strictly worse than y , while $y \in P_i(h_1^*)$. Therefore, this type has no obviously dominant action at h , which is a contradiction.

Thus, $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \emptyset$, which means there must be distinct x, y such that (i) $x, y \in PnG_i(h)$ (ii) $x \in P_i(h_1^*)$ but $x \notin P_i(h_2^*)$ and (iii) $y \in P_i(h_2^*)$ but $y \notin P_i(h_1^*)$. If

⁴⁵Since h is on path for some type such that $y \succ_i x$, it is also on path for the type \succ'_i that is the same as \succ_i , except that \succ'_i promotes payoff x to be immediately after y .

there is a type that reaches h such that $Top(\succ_i, P_i(h)) = x$, then there is also type such that reaches h such that y is ranked immediately after x ; however, this type would have no obviously dominant action at h . The same applies for any type such that $Top(\succ_i, P_i(h)) = y$. Thus, for all types that reach h , it must be that $Top(\succ_i, P_i(h)) \neq x, y$; further, by pruning, some such type is playing a strategy such that $S_i(\succ_i)(h) = a_1^*$ and x is a possible payoff. Let $Top(\succ_i, P_i(h)) = z$ for this type. The fact that $S_i(\succ_i)(h) = a_1^*$ implies that $z \in P_i(h_1^*)$ and $z \notin G_i(h)$ (if either were false, then it would not obviously dominant for this type to play a strategy such that $S_i(\succ_i)(h) = a_1^*$ and x is a possible payoff); in other words, $z \in PnG_i(h)$, and $z \in P_i(h_1^*)$. Since we just showed that $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \emptyset$, we have $z \notin P_i(h_2^*)$. Finally, consider a type \succ_i such that $Top(\succ_i, P_i(h)) = z$ and $Top(\succ_i, P_i(h) \setminus \{z\}) = y$, and note that this type has no obviously dominant action at h .⁴⁶ ■

Clinching actions are those for which i 's payoff is completely determined after following the action. Lemma A2 shows that if a game is OSP, then at every history, for all actions a with the exception of possibly one special action a^* , all payoffs that are possible following a are also guaranteeable at h ; note, however, it does *not* say that all actions but at most one are clinching actions. Indeed, it leaves open the possibility that there are several actions that can ultimately lead to multiple final payoffs for i , which can happen when different payoffs are guaranteeable for i by following different strategies in the future of the game. The next lemma shows that if this is the case, we can always construct an equivalent OSP game such that all actions except for possibly one are clinching actions.

Lemma A3. *Let Γ be an OSP game of perfect information that is pruned with respect to the obviously dominant strategy profile $(S_i(\succ_i))_{i \in \mathcal{N}}$. There exists an equivalent OSP game Γ' with perfect information such that the following hold at each h (where i is the agent called to move at h):*

- (i) *At least $|A(h)| - 1$ actions at h are clinching actions*
- (ii) *For every payoff $x \in G_i(h)$, there exists an action $a_x \in A(h)$ that clinches x for i and $i_{h'} \neq i$ for all $h' \supset (h, a_x)$.*
- (iii) *If $P_i(h) = G_i(h)$, then all actions in $A(h)$ are clinching actions and $i_{h'} \neq i$ for any $h' \supsetneq h$.*

Proof. Consider some history h of game Γ at which the mover is $i(h) = i$. By Lemma A2, all but at most one action (denoted a^*) in $A(h)$ satisfy $P_i((h, a)) \subseteq G_i(h)$; this means that any obviously dominant strategy for type \succ_i that does not choose a^* guarantees the best possible

⁴⁶Again, such a type reaches h following footnote 45. Since $z \notin G_i(h)$ and $z \in P_i(h_1^*)$, no action $a \neq a_1^*$ can obviously dominate a_1^* . However, the worst case from a_1^* is strictly worse than y (since $z \notin G_i(h)$ and $y \notin P_i(h_1^*)$), while $y \in P_i(h_2^*)$, and so a_1^* itself is also not obviously dominant.

outcome in $P_i(h)$ for type \succ_i . Define the set $\mathcal{S}_i(h) = \{S_i : S_i(h) \neq a^* \text{ and } |X(h, S_i)| = 1\}$, and notice that each $S_i \in \mathcal{S}_i(h)$ guarantees a unique payoff for i if she plays strategy S_i starting from history h , no matter what the other agents do.

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 an action a^* such that $P_i((h, a^*)) \not\subseteq G_i(h)$ in the original game (of which there can be at most one), then there is an analogous action a^* in the new game, and the subgame following a^* 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, except that at any future history $h' \supseteq h$ at which i is called on to act, all actions (and their subgames) are deleted and replaced with the subgame starting from the history (h', a') , where $a' = S_i^m(h')$ is the action that i would have played at h' in the original game had she followed strategy $S_i^m(\cdot)$. In other words, if i 's strategy was to choose some action $a \neq a^*$ 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 that she would have chosen according to $S_i^m(\cdot)$ as well. By doing so, we have created a new game in which every action (except for a^* , if it exists) at h clinches some payoff x , and further, agent i is never called upon to move again.⁴⁷

We construct strategies in Γ' that are the 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 equivalent under their respective strategy profiles.

We must also show that if a strategy profile is obviously dominant for Γ , this modified strategy profile is obviously dominant for Γ' . To see why the modified strategy profile is obviously dominant for i , note that if her obviously dominant action in the original game was part of a strategy that guarantees some payoff x , she now is able to clinch x immediately, which is clearly obviously dominant; if her obviously dominant strategy was to follow a strategy that did not guarantee some payoff x at h , this strategy must have directed i to follow a^* at h . However, in Γ' , the subgame following a^* is unchanged relative to Γ , and so i is able to perfectly replicate this strategy, which obviously dominates following any

⁴⁷More precisely, all of i 's future moves are trivial moves in which she has only one possible action; hence these histories may further be removed to create an equivalent game in which i is never called on to move again. Note that this only applies to the actions $a \neq a^*$; it is still possible for i to follow a^* at h and be called upon to make a non-trivial move again later in the game.

of the clinching actions at h in Γ' . In addition, the game is also obviously strategy-proof 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 further, every payoff that is guaranteeable at h is also clinchable at h , and i never moves again following a clinching action. This shows parts (i) and (ii). Part (iii) follows immediately from part (ii), due to greedy strategies and pruning. \square

Lemma A4. *Let Γ be an obviously strategy-proof game that is pruned with respect to the obviously dominant strategy profile $(S_i(\succ_i))_{i \in \mathcal{N}}$ and that satisfies Lemmas A1 and A3. At any h , if there exists a previously unclinicable payoff z that becomes impossible for agent i_h at h , then $C_{i_h}^{\subset}(h) \subseteq C_i(h)$*

Proof. Let h^i be a history where agent i moves such that there is a previously unclinicable payoff z that becomes impossible for i at h^i (the case for terminal histories will be dealt with next). Therefore, i moves at some strict subhistory $h \subsetneq h^i$, and the following are true:

- (a') $z \notin P_i(h^i)$
- (b') $z \in P_i(h)$ for all $h \subsetneq h^i$ such that $i_h = i$
- (c') $z \notin C_i^{\subset}(h^i)$ 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 $h \subsetneq h^i$ where she she is called to move.⁴⁸

Towards a contradiction, assume that $C_i^{\subset}(h^i) \not\subseteq C_i(h^i)$, i.e., there exists some $h' \subsetneq h^i$ such that $i_{h'} = i$ and some $x \in C_i(h')$ such that $x \notin C_i(h^i)$. Consider a type $z \succ_i x \succ_i \dots$. By the previous paragraph, at any such $\hat{h}^i \subsetneq h^i$, any obviously dominant strategy must have this type passing. Since $z \notin P_i(h^i)$ and $x \notin C_i(h^i)$, by Lemma A3, the worst case outcome from following this strategy is some y that is strictly worse than x according to \succ_i . 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

⁴⁸Since Γ is a millipede and z is not clinchable, but is possible, at any such h , it must be possible following the (unique) passing action.

at \hat{h}^i is x . This implies that passing is not obviously dominant, and thus Γ is not OSP, a contradiction.

Last, consider a terminal history \bar{h} . As above, let z be a payoff such that (a'), (b'), and (c') hold (replacing h^i with \bar{h}). Recall that for terminal histories, we define $C_i(\bar{h}) = \{y\}$ for all i , where y is the unique outcome that obtains at \bar{h} . Towards a contradiction, assume that $C_i^c(\bar{h}) \not\subseteq C_i(\bar{h})$, i.e., there exists some $h' \subsetneq \bar{h}$ such that $i_{h'} = i$ and some payoff $x \in C_i(h')$ such that $x \notin C_i(\bar{h})$. Note that (i) $z \neq y$ (because $z \notin P_i(\bar{h})$, by (a')); (ii) $z \neq x$ (by (c')); and (iii) $x \neq y$ (because $x \notin C_i(\bar{h})$). In other words, x, y, z must all be distinct payoffs for i . Consider the type $z \succ_i x \succ_i y \succ_i \dots$. By (b') and (c'), z is possible at every $h \subsetneq \bar{h}$ where i is to move, but is not clinchable at any such history. Thus, any obviously dominant strategy of type \succ_i must have agent i passing at any such history. However, at h' , i could have clinched x , and so this strategy is not obviously dominant (because y is possible from passing). Therefore, this type has no obviously dominant strategy, which is a contradiction. \square

A.4 Proof of Theorem 2

The proof follows similar steps as the proof of the analogous result for OSP in Li (2017b).

Suppose the strategic plan S_{i,h^*} is simply dominant from the perspective of $h^* \in \mathcal{H}_i$ in Γ . Then any outcome that is possible after playing S_{i,h^*} at all histories $h \in \mathcal{H}_{i,h^*}$ is weakly better than any outcome that is possible after playing $S'_i(h^*) \neq S_{i,h^*}(h^*)$ in Γ , and hence in any game Γ' that is i -indistinguishable from Γ . Hence, every strategy S_i such that $S_i(h) = S_{i,h^*}(h)$ for all $h \in \mathcal{H}_{i,h^*}$ weakly dominates any strategy S'_i such that $S'_i(h^*) \neq S_{i,h^*}(h^*)$.

Now fix h^* at which i moves and suppose that any strategy S_i such that $S_i(h) = S_{i,h^*}(h)$ for all $h \in \mathcal{H}_{i,h^*}$ weakly dominates any strategy S'_i such that $S'_i(h^*) \neq S_{i,h^*}(h^*)$ in every game Γ' that is i -indistinguishable from Γ . Consider such a Γ' in which all moves of agent i following history h^* but not in $\mathcal{H}_{i,h^*}h$ are made by Nature instead. Since, S_i weakly dominates S'_i in Γ' , we conclude that any outcome that is possible after playing S_i is weakly better than any outcome that is possible after playing S'_i in game Γ' , and hence in the i -indistinguishable game Γ . Hence, in game Γ the strategic plan S_{i,h^*} is simply dominant from the perspective of $h^* \in \mathcal{H}_i$. \blacksquare

A.5 Proof of Theorem 4

We first prove the second statement. Let Γ be a monotonic millipede game. Fix an agent i , and, for any history h^* at which i moves, let $\bar{x}_{h^*} = \text{Top}(\succ_i, P_i(h^*))$ and $\bar{y}_{h^*} = \text{Top}(\succ_i, C_i(h^*))$.

Let $\mathcal{H}_{i,h^*} = \{h \in \mathcal{H}_i(h^*) \mid h^* \subsetneq h' \subsetneq h \implies h' \notin \mathcal{H}_i\}$ be the set of one-step-foresight simple nodes. Consider the following strategic plan for any h^* :

- If $\bar{x}_{h^*} \in C_i(h^*)$, then $S_{i,h^*}(h^*) = a_{\bar{x}_{h^*}}$, where $a_{\bar{x}_{h^*}} \in A(h^*)$ is a clinching action for \bar{x}_{h^*} .
- If $\bar{x}_{h^*} \notin C_i(h^*)$, then $S_{i,h^*}(h^*) = a^*$ (i passes at h^*), and, for any other $h \in \mathcal{H}_{i,h^*}$:
 - If $P_i(h^*) \setminus C_i(h^*) \subseteq C_i(h)$, then $S_{i,h^*}(h^*) = a_{\bar{x}_{h^*}}$.
 - If $C_i(h^*) \subseteq C_i(h)$, then $S_{i,h^*}(h^*) = a_{\bar{y}_{h^*}}$.

Note that by monotonicity, at any $h \in \mathcal{H}_{i,h^*}$, one of the conditions in the last two bullet points must hold. It is straightforward to verify that this strategic plan is OSF-simple at any h^* , and thus the corresponding strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is also OSF-simple.

Now, we prove the first statement. Let Γ be a millipede game that is not monotonic, which means there exists an agent i , a history h^* at which i moves, another history $h \in \mathcal{H}_{i,h^*}$, and payoffs x and y such that $x \in (P_i(h^*) \setminus C_i(h^*)) \setminus C_i(h)$ and $y \in C_i(h^*) \setminus C_i(h)$. Notice that $x \neq y$. Without loss of generality we assume that h is the earliest history at which monotonicity is violated in this way.⁴⁹ Since both $x, y \notin C_i(h)$ by definition, there is some third payoff $z \neq x, y$ such that $z \in C_i(h)$. Let \succ_i be a type of agent i such that $Top(\succ_i, P_i(h^*)) = x$ and $Top(\succ_i, P_i(h^*) \setminus \{x\}) = y$, and let \succ'_i be a type of agent such that $Top(\succ'_i, P_i(h^*)) = x$ and $Top(\succ'_i, P_i(h^*) \setminus \{x\}) = z$. Note that for both \succ_i and \succ'_i , for any OSF-dominant plan $S_{i,h^*}(h^*) = a^*$ (because x is possible, but not clinchable at h^*).

There are two cases, depending on what is possible at h .

Case (1): $y \notin P_i(h)$. From above, $S_{i,h^*}(h^*) = a^*$. However, for any such strategic plan, the worst case outcome from the perspective of h^* is some $w \neq x, y$.⁵⁰ Since she can clinch y at h^* , and $y \succ_i w$, $S_{i,h^*}(\cdot)$ is not OSF-dominant.

Case (2): $y \in P_i(h)$. Here, there are two subcases.

Subcase (2).(i): $z \in P_i((h, a^*))$. In this case, type \succ_i has no OSF-dominant strategic plan at h^* . Again, in any such plan, we have $S_{i,h^*}(h) = a^*$. But, since $z \in P_i((h, a^*))$, for any $a \in A(h)$, the worst case from the perspective of node h^* is at best z (since both $x, y \notin C_i(h)$, by definition), which is worse than clinching y at h^* , and so $S_{i,h^*}(\cdot)$ is not OSF-dominant.

Subcase (2).(ii): $z \notin P_i((h, a^*))$. If $x \in P_i(h)$, then type \succ'_i has no OSF-dominant strategic plan at h .⁵¹ To see this, note that at h , for any strategic plan, the worst case

⁴⁹This assumption guarantees that, in a pruned game, history h^* is on path of the play for the two agent types we construct.

⁵⁰If $x \notin P_i(h)$, then this is obvious (since in this case, $y \notin P_i(h)$ either). If $x \in P_i(h)$, by definition $x \notin C_i(h)$, and so x is only possible following a pass at h , $x \in P_i((h, a^*))$. By definition of a passing action, there is some other $w \neq x$ such that $w \in P_i((h, a^*))$. Since $y \notin P_i(h)$, $w \neq y$.

⁵¹Note that here, we consider h , not h^* ; in fact, the argument actually shows something stronger, which

from passing at h is strictly worse than z (since x is possible, but not clinchable at h , and $z \notin P_i((h, a^*))$, while $z \in C_i(h)$, and so $S_{i,h}(h) \neq a^*$. However, $S_{i,h}(h)$ must equal a^* because the best case from passing is x ; a contradiction.

If $x \notin P_i(h)$, then type \succ_i has no OSF-dominant strategic plan at h^* .⁵² Once again, any such plan must have $S_{i,h^*}(h^*) = a^*$. Since $y \in P_i(h)$ but $y \notin C_i(h)$, it must be that $y \in P_i((h, a^*))$, and so there is some other $w \neq x, y$ such that $w \in P_i((h, a^*))$ (because $x \notin P_i(h)$). Therefore, from the perspective of node h^* , for any fixed plan $S_{i,h^*}(h)$, the worst case is at best w , which is strictly worse than clinching y at h^* , and thus $S_{i,h^*}(\cdot)$ is not OSF-dominant.

A.6 Proof of Lemma 1

Because of the restriction to paths of the game that are on the path of the greedy strategies for some type profile, it is sufficient to prove this lemma for pruned games. We first note the following lemma, which says that the first time an agent is called to play in a pruned SOSP game, all of her actions are associated with a unique undominated payoff, except for possibly one action, which may have two undominated payoffs. To state the lemma, define $\hat{P}_i(h) = \{x \in P_i(h) : \nexists y \in P_i(h) \text{ s.t. } y \triangleright x\}$ to be the set of possible payoffs for i at h that are undominated.

Lemma A5. *Let Γ be a pruned SOSP game. Let h_0^i be any earliest history at which agent i is called to play. Then, $|\hat{P}_i((h_0^i, a))| \leq 2$ for all $a \in A(h_0^i)$, with equality for at most one $a \in A(h_0^i)$.*

Proof of lemma. Since h_0^i is the first time i is called to move, it is on-path for all types of agent i . We first show that $|\hat{P}_i((h_0^i, a))| \leq 2$ for all $a \in A(h_0^i)$. Assume not, which means that there exists some $a \in A(h_0^i)$ such that $|\hat{P}_i((h_0^i, a))| \geq 3$. Let $x, y, z \in \hat{P}_i((h_0^i, a))$ be three distinct undominated payoffs that are possible following a . By pruning, there must be some type, \succ_i , such that action a is strongly obviously dominant at h_0^i . Without loss of generality, let $Top(\succ_i, P_i(h_0^i)) = x$.⁵³ Now, note that the worst case from action a is strictly worse than x (since y, z are possible). Again, without loss of generality, assume that $x \succ_i y \succ_i z$ (such a type exists by richness and the assumption that x, y, z are all undominated at h_0^i). For a to be strongly obviously dominant, for all other $a' \neq a$, the best case outcome for type \succ_i must

is that there is no obviously dominant action at h , and so the game in this case is not OSP, let alone OSF-dominant.

⁵²As with the previous case, the statement can actually be made stronger: there is actually no obviously dominant action at h^* .

⁵³By definition, $Top(\succ_i, P_i(h_0^i))$ must be some undominated payoff that is possible at h_0^i , and it is without loss of generality to assume it is x .

be no better than z ; in particular, this implies that for all $a' \neq a$ and all $w \in P_i((h_0^i, a'))$, $w \not\geq y$ (note that since \geq is reflexive, this includes y itself). Choose some $w \in P_i((h_0^i, a'))$ for some $a' \neq a$, and consider a type such that $Top(\succ'_i, P_i(h_0^i)) = y$ and $y \succ'_i w \succ'_i x$.⁵⁴ For this type, the worst case from a is at best x , while the best case from a' is w , so a is not strongly obviously dominant; for any $a' \neq a$, the worst case is strictly worse than y (since nothing that dominates y is possible following any $a' \neq a$), while the best case from a is y , and so no $a' \neq a$ is strongly obviously dominant either. Therefore, type \succ'_i has no strongly obviously dominant action, which is a contradiction.

Finally, we show that $|\hat{P}_i((h_0^i, a))| = 2$ for at most one $a \in A(h_0^i)$. Let a and a' be two actions such that there are two possible undominated payoffs for i following each, and, for notational purposes, let $\hat{P}_i((h_0^i, a)) = \{x, y\}$. Again, by pruning, there is some type \succ_i that selects action a as a strongly obviously dominant action, and, as above, without loss of generality, let $Top(\succ_i, P_i(h_0^i)) = x$. Since y is possible following a , in order for a to be strongly obviously dominant, the best case from any $a' \neq a$ must be no better than y ; in other words, for all $w \in P_i((h_0^i, a'))$, $w \not\geq x$ (including $w = x$ itself). Therefore, let $\hat{P}_i((h_0^i, a')) = \{w, z\}$, where, as just argued, $w, z \not\geq x$. It is also without loss of generality to assume that y and z do not dominate each other (since by supposition there are two undominated payoffs following a' , and at most one of them can be related to y via dominance). If $w \geq y$, then consider a specific type such that $Top(\succ_i, P_i(h_0^i)) = x$ and $x \succ_i z \succ_i w$ (which one again exists by richness). Since nothing that dominates x is possible following any $a' \neq a$ (including x itself), no such a' can be strongly obviously dominant for this type. Further, the worst case from a is at best y , while the best case from a' is $z \succ_i y$, and so a is also not strongly obviously dominant. Therefore, this type has no strongly obviously dominant action. If $w \not\geq y$, then consider a type such that $x \succ_i z \succ_i y \succ_i w$, and once again note that this type has no strongly obviously dominant action at h . ■

Continuing with the main proof, if a history h is payoff-relevant, then by definition $|P_i(h)| \geq 2$. Assume that there was a path of the game with two payoff-relevant histories $h_1 \subsetneq h_2$ for some agent i , and note that it is without loss of generality to assume that h_1 and h_2 are the first and second times i is called to play on the path, and that $|P_i((h_1, a))| > 1$ for some $a \in A(h_1)$. In light of the previous lemma that $|\hat{P}_i((h_1, a))| \leq 2$ for all $a \in A(h_1)$, with equality for at most one a , there are two cases.

Case (1): There exists payoff relevant histories $h_1 \subset h_2$ such that $|\hat{P}_i((h_1, a))| = 2$, where a is the unique action such that $(h_1, a) \subset h_2$.

⁵⁴If $x \geq w$ for all $w \in P_i((h_0^i, a'))$ for all $a' \neq a$, then we consider a type such that $y \succ'_i w \succ'_i z$, and make the same argument. Note that these types exist by richness and the fact that x, y, z are all mutually undominated.

In this case, action a has two undominated possible payoffs. By the previous lemma, there can only be one such action, which we will denote a_1^* . For notational purposes, define $\hat{P}_i(h_1, a_1^*) = \{x, y\}$, where x and y are both undominated payoffs. By pruning, there must be some type whose obviously dominant strategy selects a_1^* ; without loss of generality, let $Top(\succ_i, P_i(h_1, a_1^*)) = x$.

Next, we claim that for all $a' \neq a_1^*$ and all $w \in P_i((h_1, a'))$, we have $w \supseteq y$. To see this, assume that there was some such a' and w such that $w \not\supseteq y$. By the previous lemma, $w \not\supseteq x$ for all $w \in P_i(h_1)$.⁵⁵ If $y \supseteq w$, then $y \triangleright w$ (since $w \not\supseteq y$). By pruning, some type \succ'_i is selecting action a' , and it is strongly obviously dominant; however, the worst case from a' is at best w , while y is possible from a . Since $y \triangleright w$, we have $y \succ'_i w$, and so a' does not strongly obviously dominate a , which is a contradiction. If $y \not\supseteq w$, then neither y nor w dominate each other, and type $x \succ_i w \succ_i y$ has no strongly obviously dominant action at h_1 . Therefore, $w \supseteq y$ for all $w \in P_i((h_1, a'))$ and all $a' \neq a_1$. In fact, it is further the case that $w = y$; to see this, note that if there exists some $w \triangleright y$, then type $x \succ_i w \succ_i y$ again has no strongly obviously dominant action.⁵⁶ Thus, we have shown that for all $a' \neq a_1^*$, $P_i((h_1, a')) = \{y\}$.

Since h_2 is payoff relevant, there must exist some $x', y' \in \hat{P}_i(h_2)$ that are undominated.⁵⁷ Further, we claim that $y' = y$ and $x' = x$. To see the former, first note that $y \in \hat{P}_i((h_1, a_1^*))$ (i.e., y is undominated at (h_1, a_1^*)), and so we cannot have $y' \triangleright y$. If $y \triangleright y'$, then type $x \succ_i y \succ_i y'$ has no strongly obviously dominant action at h_1 . Finally, to see that $x' = x$, again note that we cannot have $x' \triangleright x$; if $x \triangleright x'$, then type $x \succ_i y \succ_i x'$ has no strongly obviously dominant action at h_1 . Thus, $\hat{P}_i(h_2) = \{x, y\}$.

Finally, note that any type that prefers $x \succ_i y$ must select action a_1^* at h_1 , and thus, h_2 is on-path; further, any type that prefers $y \succ_i x$ must select some $a' \neq a_1^*$ at h_1 ; in other words, for all types that reach h_2 , $Top(\succ_i, P_i(h_2)) = x$. Recall that $|A(h_2)| \geq 2$, and there must be at least one action such that y is a possible outcome. Label this latter action a_2 (i.e., $y \in P_i((h_2, a_2))$), and let a'_2 be some other action. By pruning, there must exist some types \succ_i and \succ'_i whose strongly obviously dominant strategies select a_2 and a'_2 , respectively. But, $Top(\succ_i, P_i(h_2)) = Top(\succ'_i, P_i(h_2)) = x$ (indeed, as just argued above, x is the top choice for all types of i that reach h_2), which implies that $x \in P_i((h_2, a_2))$ and $x \in P_i((h_2, a'_2))$. However, $y \in P_i((h_2, a_2))$, and so a_2 is not strongly obviously dominant for type \succ_i , which

⁵⁵If $w \supseteq x$ for some $w \in P_i((h_1, a'))$, then, by the lemma, $\hat{P}_i((h, a')) = \{\hat{w}\}$, and $\hat{w} \supseteq w \supseteq x$, and therefore, $\hat{w} \succ_i w \succ_i x$ for all types of agent i . But, this contradicts that a_1^* was strongly obviously dominant for type \succ_i .

⁵⁶The worst case from a_1^* is at best y , while w is possible from some a' , and so a_1^* is not strongly obviously dominant; similarly, the worst case from any $a' \neq a_1^*$ is strictly worse than x , while the best case from a_1^* is x (or otherwise, it would not be obviously dominant for the type that ranks x first to choose a_1^*).

⁵⁷If not, then all payoffs in $P_i(h_2)$ can be ordered by the dominance relation \supseteq , and, if a strongly obviously dominant action exists, all types will take the same action, and the remaining actions can be pruned.

is a contradiction.

Case (2): For all payoff-relevant histories $h_1 \subset h_2$, $|\hat{P}_i((h_1, a))| = 1$, where a is the unique action such that $(h_1, a) \subset h_2$.

Note that $|\hat{P}_i((h_1, a))| = 1$ implies that $|\hat{P}_i(h_2)| = 1$. Let $\hat{P}_i(h_2) = \{x\}$, and note that by definition, $x \succeq x'$ for all $x' \in P_i(h_2)$, which implies that for all types of agent i , $x \succsim_i x'$ for all $x' \in P_i(h_2)$. Since there are no trivial moves, $|A(h_2)| \geq 2$. Since h_2 is payoff-relevant, $|P_i(h_2)| \geq 2$, i.e., there must exist some $a'_2 \in A(h_2)$ and $x' \in P_i((h_2, a'_2))$ such that $x' \neq x$. Further, by pruning, there is some type that has a strongly obviously dominant strategy that selects a'_2 . This implies that, for any $a_2 \neq a'_2$, $x \notin P_i((h_2, a_2))$, and so $x \in P_i((h_2, a'_2))$. Again by pruning, there must be some type \succsim_i that has a strongly obviously dominant strategy that selects a_2 . But, as just argued, $x \notin P_i((h_2, a_2))$ and $x \in P_i((h_2, a'_2))$. Since all types are such that $Top(\succsim_i, P_i(h_2)) = x$, a_2 is not strongly obviously dominant, which is a contradiction.

A.7 Proof of Theorem 5

That sequential price mechanisms are SOSPP is immediate from the definition, and so we focus on proving that every SOSPP game is equivalent to a sequential price mechanism. Note first that the pruning principle continues to apply to strong obvious dominance. Also, following the same reasoning as in the proof of Proposition 1, given any SOSPP game, we can construct an equivalent SOSPP game of perfect information in which Nature moves at most once, as the first mover, and so we can focus on the deterministic subgame after any potential move by Nature. Thus, what remains to show is that every perfect-information, pruned SOSPP game in which there are no moves by Nature is equivalent to a sequential price mechanism.

Let Γ be such a game. By Lemma 1, each agent i can have at most one payoff-relevant history along any path of game Γ , and this history (if it exists) is the first time i is called to play. Consider any such history h_0^i . If there is some other history $h' \supset h_0^i$ at which i is called to play, then history h' must be payoff-irrelevant for i ; in other words, there is some payoff x such that $P_i((h', a')) = \{x\}$ for all $a' \in A(h')$. Using the same technique as in the proof of Theorem 1, we can construct an equivalent game Γ' in which at history h_0^i , i is asked to also choose her actions for all successor histories $h' \supset h_0^i$ at which she might be called to play, and then is not called to play again after h_0^i (see the proof of Theorem 1 for a more formal description of this procedure). Since all of these future histories were payoff-irrelevant for i , the new game continues to be strongly obviously dominant for i . Strong obvious dominance is also preserved for all $j \neq i$, since having i make all of her choices earlier only shrinks the set of possible outcomes any time j is called to move, and thus, if

some action was strongly obviously dominant in the old game, the analogous action(s) will be strongly obviously dominant in the new game. Repeating this for every agent and every history, we have constructed a SOSPP game Γ' that is equivalent to Γ and in which each agent is called to move at most once along any path of play.

We claim that Γ' is a sequential price mechanism. Assume not, and let h be an earliest history where the definition of a sequential price mechanism is violated. Since Γ' is not a sequential price mechanism, there must be some payoff $x \in P_i(h)$ that i cannot clinch at h . Note that it is without loss of generality to assume that there exists such an unclinched x that is not dominated, i.e., $x \in \hat{P}_i(h)$.⁵⁸ Since x is not clinchable, for any action $a \in A(h)$ such that $x \in P_i((h, a))$, there is some other $y \in P_i((h, a))$.

Case (1): $|P_i(h)| \geq 3$ and there exists a $y \in P_i((h, a))$ such that x and y do not dominate each other.

By Lemma A5 (in the proof of Lemma 1), a is the unique action such that $|\hat{P}_i((h, a))| = \{x, y\}$, and, for any other $a' \neq a$, let $\hat{P}_i((h, a')) = \{w'\}$. We first claim that for any $a' \neq a$, $\hat{P}_i((h, a')) = \{y\}$.

Assume not, i.e., there exists some $a' \neq a$ and $w' \neq y$ such that $\hat{P}_i((h, a')) = \{w'\}$. First, since x is not clinchable, any type such that $Top(\succ_i, P_i(h)) = x$ must select a , and $x \notin P_i((h, a'))$ for any $a' \neq a$.⁵⁹ Now, if $x \triangleright w'$, then type $x \succ_i w' \succ_i y$ has no strongly obviously dominant action at h ; therefore, $x \not\triangleright w'$. If $y \triangleright w'$, then $y \notin P_i((h, a'))$ (since by assumption w' is undominated at (h, a')); however, if this is the case, then it is not strongly obviously dominant for any type to select a' (since $y \in P_i((h, a))$), and it can be pruned. Therefore, $y \not\triangleright w'$. If $w' \triangleright y$, then, once again, type $x \succ_i w' \succ_i y$ has no strongly obviously dominant action at h . Therefore, $w' \not\triangleright y$. Thus, the only remaining possibility is that $x, y, w' \in \hat{P}_i(h)$, i.e., x, y, w' are all mutually undominated payoffs at h . But then, type $x \succ_i w' \succ_i y$ has no strongly obviously dominant action at h . Therefore, $\hat{P}_i((h, a')) = \{y\}$ for all $a' \neq a$.

We also claim further that $P_i((h, a')) = \{y\}$ for all $a' \neq a$; indeed, if this were not the case, then there is some a' and some $w' \in P_i((h, a'))$ such that $y \triangleright w'$. By pruning, some type \succ'_i must be selecting action a' . However, the worst case from a' (for all types) is at best w' , while y is possible following a , and so a' is not strongly obviously dominant for type \succ'_i . Therefore, $P_i((h, a')) = \{y\}$ for all $a' \neq a$.

Let $z \neq x, y$ be some third payoff that is possible at h . In light of the previous paragraph, $z \in P_i((h, a))$, and $z \notin P_i((h, a'))$ for all other $a' \neq a$. Finally, note that type $x \succ_i y \succ_i z$

⁵⁸If all $x' \in \hat{P}_i(h)$ are clinchable at h , then all types will be able to take an action that clinches their top possible payoff, and any other action can be pruned.

⁵⁹Recall that x is not dominated at h , so such a type does indeed exist.

has no strongly obviously dominant action at h . (Note that since $\hat{P}_i(h) = \{x, y\}$, z must be dominated by one of x or y , and so by our richness assumption, such a type exists.)

Case (2): $|P_i(h)| \geq 3$ and for all other $y \in P_i((h, a))$, $x \triangleright y$.

Since i cannot clinch x at h , we have that, for all other a' and $w \in P_i((h, a'))$, $x \triangleright y \succeq w$.⁶⁰ By pruning, some type \succ'_i must be choosing action a' ; however, the previous sentence implies that a' is not strongly obviously dominant for this type, which is a contradiction.

Case (3): $|P_i(h)| = 2$.⁶¹

Let $P_i(h) = \{x, y\}$. Given the definition of a sequential price mechanism, the only case we need to rule out is that neither x nor y is clinchable, i.e., there are at least two actions in $A(h)$, and, for all $a \in A(h)$, $P_i((h, a)) = \{x, y\}$. At least one of $x \succ_i y$ or $y \succ_i x$ must hold for some type at h ; however, it is simple to see that no matter which is true, this will not have a strongly obviously dominant action. ■

A.8 Proof of Theorem 6

A.8.1 Roles and “Symmetry to Symmetrization” Reduction Lemma

For clarity of the exposition, it is convenient to sometimes distinguish between an agent i moving at some set of histories and a “role” moving at these same histories.⁶² Formally, we create a copy \mathcal{R} of the set of agents \mathcal{N} . Given a perfect-information mechanism (Γ, S) we create a copy of the game Γ as a game between these roles—treated as agents—and we create a copy of the strategy profile S as strategies of these roles.⁶³ With some abuse of notation we refer to the copy of (Γ, S) by the same symbols. For a game Γ , the function $\rho : \mathcal{H} \rightarrow \mathcal{R}$ maps each history h to the role $\rho(h)$ that moves at this history.

We use the role copy of (Γ, S) to create mechanisms $(\Gamma_\sigma, S_\sigma)$ that differ only in the mapping of agents (and their preferences) to the roles. The preferences of the roles are determined by the preferences of the original agents and a bijection $\sigma : \mathcal{R} \rightarrow \mathcal{N}$. We call this bijection a **role assignment** function, and we denote by Σ the space of all role assignment functions. We define Γ_σ as the extensive-form game with the same game tree as Γ and such that at each non-terminal history h , the agent called to move is $\sigma(\rho(h))$; at each terminal history in Γ_σ the payoff of agent $\sigma(i)$ is the same as the payoff of i at the corresponding

⁶⁰Note that in particular, this implies that $x \succeq w$ for all $w \in P_i(h)$. Indeed, if there were some $w \in P_i(h)$ such that w and x did not dominate each other, then type $x \succ_i w \succ_i y$ has no strongly obviously dominant action at h . (Recall also that x is undominated at h , and so there is no $w \triangleright x$, either.)

⁶¹If $|P_i(h)| = 1$, the argument is trivial.

⁶²Our construction of roles in general extensive-form games extends the role concept from Carroll (2014), who studied them in the (static) context of Pápai (2000)’s hierarchical exchange mechanisms.

⁶³While in the main text we denote a profile of strategies in a mechanism as $S_{\mathcal{N}}$, here we just write S to avoid notational clutter.

history in Γ . The strategy of agent $\sigma(i)$ is the same as the strategy of agent i in the original game Γ . There are $|\Sigma| = N!$ possible mechanisms $(\Gamma_\sigma, S_\sigma)$; we call them the permuted mechanisms.

We further define the **symmetrized mechanism** (Γ^*, S^*) to be the following random mechanism: first, Nature chooses a role assignment function σ uniformly at random from the set of all possible role assignment functions, and then, the agents play Γ_σ with strategies S_σ .⁶⁴ To formally ensure that the symmetrizations of a millipede is a millipede, we assume that Nature draws the role assignment σ and the path in the subgame Γ_σ in the same move.

The following lemma shows that it is sufficient to prove Theorem 6 for symmetrized mechanisms.

Lemma A6. *Suppose that every symmetrization of a deterministic OSP and Pareto-efficient perfect-information mechanism is equivalent to Random Priority. Then, every symmetric, OSP and Pareto-efficient mechanism is equivalent to Random Priority.*

Proof. Take a symmetric, OSP, and Pareto-efficient mechanism (Γ, S) . By Lemma A1, we can assume that (Γ, S) has perfect information and that Nature moves only at the beginning of the game. Because (Γ, S) is symmetric the symmetrized mechanism (Γ^*, S^*) is equivalent to (Γ, S) . Furthermore, (Γ^*, S^*) is a lottery over symmetrizations of each deterministic perfect-information continuation game Γ' after Nature's move in (Γ, S) . The mechanism given by game Γ' , together with the strategy profile induced from Γ , is OSP and Pareto efficient, and hence by the assumption of the lemma it is equivalent to Random Priority. Because every lottery over Random Priority lotteries is still equivalent to Random Priority, the lemma obtains. ■

In light of the above lemma, in the sequel we focus on symmetrized mechanisms.

A.8.2 Plan of the Remainder of the Proof

The remainder of the proof builds on the bijective argument used by Abdulkadiroğlu and Sönmez (1998) to show the equivalence of Random Priority and the Core from Random Endowments (see also Pathak and Sethuraman, 2011 and Carroll, 2014). Throughout, we fix the profile of preferences $\succ_{\mathcal{N}}$. Given any Γ that is OSP and Pareto efficient, we construct a bijection $f : \Sigma \rightarrow Ord$ that associates to each role assignment function $\sigma \in \Sigma$ a total linear order of the agents $f_\sigma \in Ord$ with the property that game Γ_σ results in the same final

⁶⁴While this construction implies that different agents play the same strategies in the same role, our arguments only rely on the weaker assumption that an agent's strategy $S_{\sigma,i}(\succ_i)$ depends only on her own preferences and her role assignment, and not on the roles assigned to other agents. In other words, in any two subgames Γ_A and Γ_B following Nature's selection of role assignments σ_A and σ_B , if $\sigma_A^{-1}(i) = \sigma_B^{-1}(i) = r_n$, then $S_i^*(\succ_i)(h_A) = S_i^*(\succ_i)(h_B)$ for any equivalent histories h_A and h_B in these two games.

allocation (matching) μ as a serial dictatorship where the first agent called to play is $f_\sigma(1)$, the second agent called to play is $f_\sigma(2)$, etc. We then show that the mapping we constructed is a bijection, which proves Theorem 6.

A.8.3 Efficient Millipedes

With slight abuse of notation, in the remainder of the proof, we use \mathcal{X} to denote the set of objects to be allocated (rather than global outcomes), and use x, y, z , etc. to refer to objects from \mathcal{X} .

Given any OSP game Γ , we can use Lemma A3 to construct an equivalent millipede game which has the following properties:

1. At each history h , there is at most one passing action in $A(h)$; this action, if it exists, is denoted $a^* \in A(h)$.
2. For every $x \in G_i(h)$, there exists a clinching action $a_x \in A(h)$ that clinches x for i .
3. As soon as an agent's top still-possible object is guaranteeable at a history h , she clinches this object at h (that is, agents follow greedy strategies).
4. If $i_h = i$ and $P_i(h) = G_i(h)$, then i clinches her payoff immediately at h , and is not called to move at any $h' \supseteq h$.

Agent i is **active** at h if she has been previously called to play at some $h' \subseteq h$, and further has not yet clinched an object at h . Let $\mathcal{A}(h)$ denote the set of active agents at h .

In constructing the bijection f , we make use of the concept of a lurker introduced by Bade and Gonczarowski (2017, hereafter BG).⁶⁵ Informally, a lurker is an agent who has been offered to clinch all objects that are possible for him except for exactly one, which he is said to “lurk”. If an agent lurks some object x , then the mechanism can infer that x is his favorite (still available) object, and so it is possible to exclude x from other agents without violating Pareto efficiency. The role of lurkers is to allow more than two agents to be active at any given point of the game; while there can be an arbitrary number of lurkers, at any point, at most two active agents are non-lurkers.

To formally define lurker, recall that $C_i^{\subseteq}(h) = \{x : x \in C_i(h') \text{ for some } h' \subseteq h\}$ is the objects agent i has been offered to clinch at some subhistory of h and $C_i^{\subset}(h) = \{x : x \in C_i(h') \text{ for some } h' \subsetneq h\}$ is the objects agent i has been offered to clinch at some strict subhistory of h . We consider a history h and an active agent i who has moved at a strict

⁶⁵They focus on understanding which OSP mechanisms are Pareto efficient. While in this proof we build on their insights, in turn their analysis follows our 2016 characterization of OSP mechanisms through millipede games as well as our analysis of SOSp and efficient mechanisms.

subhistory of h . Let $h' \subsetneq h$ be the maximal strict subhistory such that $i_{h'} = i$. Agent i is said to be a **lurker** for object x at h if (i) $x \in P_i(h')$, (ii) $C_i^{\subseteq}(h') = P_i(h') \setminus \{x\}$ and (iii) $x \notin C_j^{\subseteq}(h')$ for any other active $j \neq i$ that was not already a lurker prior to h' .⁶⁶

At any h , we partition the set of active agents as $\mathcal{A}(h) = \mathcal{L}(h) \cup \bar{\mathcal{L}}(h)$, where $\mathcal{L}(h) = \{\ell_1^h, \dots, \ell_m^h\}$ is the set of lurkers and $\bar{\mathcal{L}}(h)$ is the set of active non-lurkers. With some abuse of notation, we let $\mathcal{X}(h)$ denote the set of still-available (unclinched) objects at h (rather than outcomes), and partition this set as $\mathcal{X}(h) = \mathcal{X}^{\mathcal{L}}(h) \cup \bar{\mathcal{X}}^{\mathcal{L}}(h)$, where $\mathcal{X}^{\mathcal{L}}(h) = \{x_1^h, \dots, x_{\lambda(h)}^h\}$ is the set of lurked objects and $\bar{\mathcal{X}}^{\mathcal{L}}(h) = \mathcal{X}(h) \setminus \mathcal{X}^{\mathcal{L}}(h)$ is the set of unlurked objects at h . Each ℓ_m^h has a unique object that she lurks, x_m^h , and the sets are ordered such that if $m' < m$, then lurker $\ell_{m'}^h$ is “older” than lurker ℓ_m^h , in the sense that $\ell_{m'}^h$ first became a lurker for $x_{m'}^h$ at a strict subhistory of the history at which ℓ_m^h became a lurker for x_m^h .⁶⁷

In a millipede game, at any history, there is a set of clinching actions and (possibly) one passing action. Along any game path, agents engage in a sequence of passes, and the set of lurkers/lurked objects continues to grow, until eventually, we reach a history h where some agent i clinches some object x . BG show that at most two active agents are non-lurkers at any point (see Lemma A13 below). When i clinches at h , this initiates a chain of clinching among the active agents that proceeds as follows:

- If $x \in \bar{\mathcal{X}}^{\mathcal{L}}(h)$, each lurker $\ell_m^h \in \mathcal{L}(h)$ is immediately assigned to her lurked object, x_m^h .
- If $x = x_{m_1}^h$ for some lurked $x_{m_1}^h \in \mathcal{X}^{\mathcal{L}}(h)$, then all “older” lurkers $\ell_{m'}^h$ for $m' < m_1$ receive their lurked objects $x_{m'}^h$; since lurker $\ell_{m_1}^h$ ’s lurked object was taken by i , she is offered to clinch anything from the remaining set of unclinched objects, $\mathcal{X}(h) \setminus \{x_1^h, \dots, x_{m_1}^h\}$.
- If $\ell_{m_1}^h$ takes an unlurked object, then all remaining lurkers get their lurked objects; if $\ell_{m_1}^h$ chooses a lurked object $x_{m_2}^h$ for some $m_2 > m_1$, then all “older” unmatched lurkers ($\ell_{m'}^h$ for $m_1 < m' < m_2$) get their lurked objects. Lurker $\ell_{m_2}^h$ gets to choose from $\mathcal{X}(h) \setminus \{x_1^h, \dots, x_{m_2}^h\}$, etc.
- This process is repeated until some lurker ℓ_m^h chooses an unlurked object, y , at which point all remaining unassigned lurkers are assigned to their lurked objects.
- Finally, if $y \in C_j^{\subseteq}(h)$ for the other active non-lurker $j \in \bar{\mathcal{L}}(h) \setminus \{i\}$,⁶⁸ then j is offered to clinch anything from what remains, $\bar{\mathcal{X}}^{\mathcal{L}}(h) \setminus \{y\}$.

⁶⁶This definition of a lurker modifies Definition E.9 of BG, who do not impose (iii) and impose instead the requirement that $P_i(h) \neq G_i(h)$; when restricted to millipede games that satisfy properties 1-4, the definitions are equivalent.

⁶⁷That this entire construction is well-defined follows from a series of lemmas in the appendix of Bade and Gonczarowski (2017). These lemmas will also be useful in our proof, and so for ease of reference, we present in them Section A.8.6 below.

⁶⁸Such an agent may or may not exist, but if they do, they are unique.

Notice that at the end of the above chain of lurker assignments that was initiated at h , all but at most one active agent in $\mathcal{A}(h)$ has clinched and are thus no longer active. If all active agents have been assigned, then the continuation game is just a smaller Pareto efficient millipede game on the remaining unmatched agents and objects, which proceeds in the same way. If there is one active agent left, say j , then this continuation game begins with agent j carrying over anything that she has been previously offered to clinch, $C_j^{\subseteq}(h)$.

A.8.4 Constructing the bijection

We now construct the bijection from role assignment functions into serial dictatorship orderings. We start by providing an ordering algorithm that, for a give game Γ and fixed preference profile/strategy, follows the path of the game from the root node h_0 to the terminal node \bar{h} and outputs a partial ordering of the agents, denoted \triangleright .⁶⁹ This ordering is only partial because agents may “tie”. Each role assignment function $\sigma \in \Sigma$ induces a game Γ_σ and an associated partial ordering, \triangleright_σ , via our ordering algorithm. Running the algorithm on all $N!$ role assignment functions gives $N!$ partial orderings. We will then argue that it is possible to “break ties” consistently in such a way that we recover a bijection $f : \Sigma \rightarrow Ord$ such that for each σ , a serial dictatorship run under ordering f_σ results in the same allocation as game Γ_σ .

The intuitive idea behind constructing the partial order \triangleright is as follows. We start by finding the first agent to clinch some object x (after a series of passes) at some history h . This induces a chain of assignments of the active agents $\mathcal{A}(h)$ as described above. We create \triangleright by ordering agents who receive lurked objects in order of the “age” of the lurked object they receive, i.e., the agent who receives the “oldest” lurked object is ordered first, etc. After this is done, there are at most 2 active agents who have yet to be ordered, one of whom has clinched an unlurked object, say y ; if y was previously offered to the remaining active agent, then we add both remaining agents to the order without distinguishing between them, i.e., these two agents tie; if y was not previously offered to the other remaining active agent, then we just add the agent who clinched y , and the other active agent (if they exist) carries over their “endowment” (the set $C_j^{\subseteq}(h)$) to the next stage. After clearing this first segment of agents, we continue along the game path and find the first unordered agent to clinch an object, and repeat.

Ordering Algorithm. Consider any game path from the root node h_0 to a terminal node \bar{h} , which is associated with a unique allocation of objects to agents. Each step k of the algorithm below produces a partial ordering $\tilde{\triangleright}^k$ on the set of agents who are processed

⁶⁹Despite the use of the same symbol (\triangleright), the partial ordering here is unrelated to the dominance relation used to define partitions introduced in Section 2.

in step k . At the end of the final step K , we concatenate the K components to produce \triangleright , the final partial ordering on the set of all agents \mathcal{N} .

Step 1 Find the first object to be clinched along the game path, say x^1 at history h^1 by agent i^1 .⁷⁰ Let $\mathcal{L}(h^1) = \{\ell_1, \dots, \ell_{\lambda(h^1)}\}$ be the set of lurkers, and $\mathcal{X}^{\mathcal{L}}(h^1) = \{x_1, \dots, x_{\lambda(h^1)}\}$ be the set of lurked objects (note that these sets may be empty, in which case skip immediately to step 1.2 below).

1. Let i_{x_1} be the agent who ultimately receives x_1 , i_{x_2} be the agent who ultimately receives x_2 , up to $i_{x_{\lambda(h^1)}}$ (note that i_{x_k} is not necessarily the agent who lurks x_k at h^1 , but the agent who ultimately receives x_k at the allocation associated with \bar{h}).
2. Let $j \in \mathcal{L}(h^1) \cup \{i^1\}$ be the unique agent that is not one of the i_{x_k} from step 1.1. By construction, j clinches some unlurked object $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h^1)$. In addition, there may be one other active agent $j' \in \mathcal{A}(h^1) \setminus (\mathcal{L}(h^1) \cup \{i^1\})$.

(a) If such a j' exists and $y \in C_{j'}^{\subset}(h^1)$, then define $\tilde{\succ}^1$ as:

$$i_{x_1} \tilde{\succ}^1 i_{x_2} \tilde{\succ}^1 \dots \tilde{\succ}^1 i_{x_{\lambda(h^1)}} \tilde{\succ}^1 \{j, j'\}$$

(b) Otherwise, define $\tilde{\succ}^1$ as

$$i_{x_1} \tilde{\succ}^1 i_{x_2} \tilde{\succ}^1 \dots \tilde{\succ}^1 i_{x_{\lambda(h^1)}} \tilde{\succ}^1 j$$

In particular, we do not yet order agent j' .

Step k Find the first object to be clinched along the game path by an agent that has not yet been ordered, say x^k at history h^k by agent i^k . Let $\mathcal{L}(h^k) = \{\ell_1, \dots, \ell_{\lambda(h^k)}\}$ be the set of lurkers, and $\mathcal{X}^{\mathcal{L}}(h^k) = \{x_1, \dots, x_{\lambda(h^k)}\}$ be the set of lurked objects, and carry out a procedure analogous to that from step 1 to produce the step k order $\tilde{\succ}^k$.⁷¹

This produces a collection of partial orderings $(\tilde{\succ}^1, \dots, \tilde{\succ}^K)$, where each $\tilde{\succ}^k$ is a partial order on the agents processed in step k . We then create the final \triangleright in the natural way: for any two agents i, j who were processed in the same step k , $i \triangleright j$ if and only if $i \tilde{\succ}^k j$. For any two agents i, j processed in different steps $k < k'$, respectively, we order $i \triangleright j$.

⁷⁰That is, $i_{h^1} = i^1$, and i^1 selects a clinching action $a_{x^1} \in A(h^1)$ that clinches x^1 . Also, by Lemma A15, $i^1 \notin \mathcal{L}(h^1)$.

⁷¹At the end of step $k-1$, there is at most one active agent $j' \in \mathcal{A}(h^{k-1})$ that was not ordered in step $k-1$. This agent j' , if she exists, is the active non-lurker other than the non-lurker i^{k-1} that clinched at h^{k-1} to initiate the step $k-1$ assignments. Thus, after the step $k-1$ assignments are all made, we are left with a subgame where agent j carries over her previous endowment, $C_j^{\subset}(h^{k-1})$. This subgame is again a Pareto efficient millipede game, and so the same structure as the original game, but among only the unmatched agents and unclinched objects after step $k-1$. At the ‘‘root node’’ of this subgame, h_0^k , agent j is offered to clinch $C_j(h_0^k) \supseteq C_j^{\subset}(h^{k-1})$. All of the structure and arguments from the previous steps are then repeated.

Remark. The output of the ordering algorithm is a partial order, \triangleright , on \mathcal{N} , the set of agents. If there are two agents j and j' such that $j \not\triangleright j'$ and $j' \not\triangleright j$, then we say j and j' **tie** under \triangleright . Note that by construction, all ties are of size at most 2, and agents can only tie if they are processed in the same step of the algorithm.

A.8.5 Completing the proof

We complete the proof Theorem 6 using three key lemmas relating to properties of the partial orders produced by the ordering algorithm applied to games with different role assignments. The proofs of these lemmas are somewhat involved, and so to streamline the presentation of the main argument, we relegate them to the following subsections.

Take a role assignment function σ , corresponding game Γ_σ , and the partial ordering \triangleright_σ that results from applying the ordering algorithm to Γ_σ . Let f be a total ordering of the agents, where $f(1) = i$ is the first agent, $f(2) = j$ is the second agent, etc. We say that f is **consistent** with \triangleright_σ if, for all j, j' : $j \triangleright_\sigma j'$ implies $f^{-1}(j) < f^{-1}(j')$. In other words, given some partial ordering \triangleright_σ , total order f is consistent if there is some possible way to break the ties in \triangleright_σ that delivers f .

Lemma A7. *For any total order f consistent with \triangleright_σ , a serial dictatorship under agent ordering f results in the same final allocation as Γ_σ .*

For the next lemma, let h_A^k be the history that initiates step k of the ordering algorithm when it is applied to game Γ_A . For instance, $h_A^1 = (h_\emptyset, a^*, \dots, a^*)$ is a history following a sequence of passes such that agent $i_{h_A^1}$ moves at h_A^1 and is the first agent to clinch in the game. This induces a chain of assignments of the agents in $\mathcal{L}(h_A^1) \cup \{i_{h_A^1}\}$, plus possibly one other active non-lurker at h_A^1 , as described above. History $h_A^2 \supseteq h_A^1$ is then the next time along the game path that an agent who was not ordered in step 1 of the ordering algorithm clinches an object, etc.

Lemma A8. *Let σ_A, σ_B be two role assignment functions, Γ_A and Γ_B their associated games, and $(\tilde{\triangleright}_A^1, \dots, \tilde{\triangleright}_A^{K_A})$ and $(\tilde{\triangleright}_B^1, \dots, \tilde{\triangleright}_B^{K_B})$ the respective partial orderings produced by each step of the ordering algorithm. For all k , if $\tilde{\triangleright}_A^k = \tilde{\triangleright}_B^k$, then $h_A^k = h_B^k$, and further, $\sigma_A^{-1}(i) = \sigma_B^{-1}(i)$ for all agents i that are ordered in step k of algorithm.*

In particular, Lemma A8 implies the following corollary.

Corollary A1. *If $\triangleright_A = \triangleright_B$, then $\sigma_A = \sigma_B$.*

Lemma A9. *There are no three role assignment functions σ_A, σ_B and σ_C such that the resulting partial orders $\triangleright_A, \triangleright_B$ and \triangleright_C take the form:⁷²*

$$\begin{aligned} i_1 \triangleright_A \cdots \triangleright_A i_n \triangleright_A \{i, j\} \cdots \\ i_1 \triangleright_B \cdots \triangleright_B i_n \triangleright_B i \triangleright_B j \cdots \\ i_1 \triangleright_C \cdots \triangleright_C i_n \triangleright_C j \triangleright_C i \cdots . \end{aligned}$$

By Corollary A1 of Lemma A8, the mapping from role assignments σ to partial orders \triangleright_σ generated by the ordering algorithm is an injection. Lemma A9 shows that we can break all the ties—recursively, coding step by coding step—creating from each \triangleright_σ a consistent total order f_σ in a way that preserves the injectivity. In this way we obtain an injection from role assignments σ to serial dictatorships with orders f_σ . Because in this injection the domain of role assignments σ and the range of serial dictatorship orderings f_σ are finite and have equal size, this injection is actually a bijection. It remains to check that the millipede Γ_σ generates the same allocation as the serial dictatorship with ordering f_σ . This is implied by Lemma A7 because, by definition, each complete order f_σ generated by the tie-breaking in partial order \triangleright_σ is consistent with \triangleright_σ . ■

A.8.6 Preliminary Results for the Proofs of the Key Lemmas

Before proving the core Lemmas A7, A8, and A9, we state several preliminary lemmas we will use. Lemmas A10-A14 are due to BG. Note that the versions presented here are simplifications of the corresponding lemmas in BG to apply to millipede games that satisfy the properties of Lemma A3. Lemmas A15 and A16 are new.

Throughout, we fix a Pareto efficient millipede game Γ that satisfies properties (i)-(iii) of Lemma A3.

Lemma A10. *(BG Lemma E.11) If an agent has not yet clinched an object at a history h , then $\bar{\mathcal{X}}^\mathcal{L}(h) \subseteq P_i(h) \cup C_i^c(h)$. If $i \in \mathcal{L}(h)$, then $\bar{\mathcal{X}}^\mathcal{L}(h) \subseteq C_i^c(h)$.*

Lemma A11. *(BG Lemma E.14) If $i \in \bar{\mathcal{L}}(h)$ and $x_\ell \in C_i^c(h)$ for some $x_\ell \in \mathcal{X}^\mathcal{L}(h)$, then $i_h = i$, $P_i(h) = G_i(h) = C_i(h)$, and there is no passing action a^* in $A(h)$.*

Lemma A12. *(BG Lemma E.16) Let $\mathcal{L}(h) = \{\ell_1^h, \dots, \ell_{\lambda(h)}^h\}$ be the set of lurkers at h and $\mathcal{X}^\mathcal{L}(h) = \{x_1^h, \dots, x_{\lambda(h)}^h\}$, with ℓ_1^h lurking x_1^h , ℓ_2^h lurking x_2^h , etc., where $m < m'$ if and only if ℓ_m^h became a lurker at a strict subhistory of the history at which $\ell_{m'}^h$ became a lurker. Then,*

⁷²What is meant here is that σ_A, σ_B , and σ_C restricted to the agents $\{i_1, \dots, i_n, i, j\}$ are all distinct role assignment functions that produce partial orderings $\triangleright_A, \triangleright_B$, and \triangleright_C that begin by ordering agents i_1, \dots, i_n in the exact same way (possibly with ties), and continue by ordering agents i, j in the manner specified.

1. $x_1^h, \dots, x_{\lambda(h)}^h$ are all distinct objects.
2. For all $m = 1, \dots, \lambda(h)$, $P_{\ell_m^h}(h) = \mathcal{X}(h) \setminus \{x_1^h, \dots, x_{m-1}^h\}$.

Lemma A13. (BG Lemma E.19) For all h , $|\bar{\mathcal{L}}(h)| \leq 2$.

Lemma A14. (BG Lemma E.18, E.20) Let h be a history with lurked objects and let $i_{h'} = t$ be the agent who moves at the maximal superhistory of the form $h' = (h, a^*, \dots, a^*)$. Then:

- (i) Agent t is not a lurker at h .
- (ii) If $i_h \neq t$, then $C_{i_h}(h) \cap C_t^{\subseteq}(h) = \emptyset$.
- (iii) If $x_\ell \in P_j(h)$ for some non-lurker j and lurked object $x_\ell \in \mathcal{X}^{\mathcal{L}}(h)$, then $j = t$ and $C_j^{\subseteq}(h') = \mathcal{X}(h)$.

The agent t who moves at h' is called the **terminator**.

We also prove the following additional lemmas.

Lemma A15. Let h be a history such that $\mathcal{L}(h) \neq \emptyset$. For any superhistory h' of the form $h' = (h, a^*, a^*, \dots, a^*)$, we have $i_{h'} \notin \mathcal{L}(h)$.

Lemma A15 has the following key implication: let h be a history with lurkers $\mathcal{L}(h)$, and $h' = (h, a^*, \dots, a^*, a_x)$ be a superhistory such that x is the next object to be clinched (with possibly agents passing in the mean time). Then, the agent that clinches x is not a lurker.

Proof. Let $\mathcal{L}(h) = \{\ell_1^h, \dots, \ell_{\lambda(h)}^h\}$ be the set of lurkers at h and $\mathcal{X}^{\mathcal{L}}(h) = \{x_1^h, \dots, x_{\lambda(h)}^h\}$ the set of lurked objects. Assume that the statement was false, and let h' be the smallest superhistory of h such that $i_{h'} = \ell_m^h$ for a lurker ℓ_m^h (that is, $i_{h''} \notin \mathcal{L}(h)$ for all $h \subseteq h'' \subsetneq h'$). Note first that, for any h'' such that $h \subseteq h'' \subsetneq h'$, $i_{h''} = j \in \bar{\mathcal{L}}(h)$, and if there exists some lurked $x_m^h \in C_j^{\subseteq}(h'')$, by BG Lemma A11, there is no passing action at h'' , which is a contradiction. Therefore, any clinching action $a_y \in A(h'')$ clinches some $y \in \mathcal{X}(h) \setminus \mathcal{X}^{\mathcal{L}}(h)$, and for all terminal histories $\bar{h} \supset (h'', a_y)$, each lurker $\ell_m^h \in \mathcal{L}(h)$ receives his lurked object x_m^h . Finally, consider history h' . By BG Lemma A12, for each $\ell_m^h \in \mathcal{L}(h)$, $P_{\ell_m^h}(h') = P_{\ell_m^h}(h) = \mathcal{X}(h) \setminus \{x_1^h, \dots, x_{m-1}^h\}$ (note that h' is reached from h via a series of passes, and so $\mathcal{X}(h) = \mathcal{X}(h')$), and $Top(\succ_{\ell_m^h}, P_{\ell_m^h}(h')) = x_m^h$ for all types $\succ_{\ell_m^h}$ such that h' is on the path of play. Therefore, by pruning and greedy strategies, at h' , there is no clinching action a_x for any $x \in P_{\ell_m^h}(h') \setminus \{x_m^h\}$. Thus, the only possibility is that every action $a \in A(h')$ clinches x_m^h .⁷³ This then implies that ℓ_m^h gets x_m^h at all terminal $\bar{h} \supset h'$. Combining this with the previous statement that ℓ_m^h gets x_m^h for all terminal $\bar{h} \supset (h'', a_y)$ for any $h \subseteq h'' \subsetneq h'$ and

⁷³Note that there cannot be a passing action either: if there were, then, since every history is non-trivial, there must be another action. But, as just argued, there can be no clinching actions for any other $x \neq x_m^h$, and thus there must be a clinching action for x_m^h , and the passing action would be pruned.

clinchng action $a_y \in A(h'')$, we conclude that ℓ_m^h gets x_m^h for all terminal $\bar{h} \supset h$, i.e., ℓ_m^h has already clinched his object x_m^h at h . Thus, by definition of a millipede game, $i_{h'} \neq \ell_m^h$, which is a contradiction. ■

Lemma A16. *Let h be a history such that $\mathcal{L}(h) \neq \emptyset$ and $\mathcal{A}(h) = \mathcal{L}(h) \cup \{i, j\}$, where i and j are active non-lurkers at h , and let $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h)$ be an unlurked object at h . Further, assume that (i) $i_h = i$ (ii) $y \in C_i(h) \cap C_j^{\subseteq}(h)$, and define $\bar{x} = \text{Top}(\succ_j, \bar{\mathcal{X}}^{\mathcal{L}}(h))$. Then, $\bar{x} \succ_j y$.*

Proof. Let $h' \subsetneq h$ be the largest subhistory such that $y \in C_j(h')$, and note that for this history, $P_j(h') = P_j(h)$.⁷⁴ By construction, j passed at h' when she was offered to clinch y . If $P_j(h) \cap \mathcal{X}^{\mathcal{L}}(h) = \emptyset$, then, by BG Lemma A14, j is the terminator (i.e., $j = t$), and so by that same lemma, $C_i(h) \cap C_j^{\subseteq}(h) = \emptyset$, which is a contradiction (since $y \in C_i(h) \cap C_j^{\subseteq}(h)$). Therefore, $P_j(h) = \bar{\mathcal{X}}^{\mathcal{L}}(h)$,⁷⁵ and so $P_j(h') = \bar{\mathcal{X}}^{\mathcal{L}}(h)$ as well. Since j passes at h' and $y \in C_j(h')$, $\text{Top}(\succ_j, P_j(h')) \succ_j y$. Since $P_j(h') = \bar{\mathcal{X}}^{\mathcal{L}}(h)$, $\text{Top}(\succ_j, P_j(h')) = \text{Top}(\succ_j, \bar{\mathcal{X}}^{\mathcal{L}}(h)) = \bar{x} \succ_j y$, as required. ■

A.8.7 Proofs of the Key Lemmas A7, A8, and A9

In the proofs that follow, we will often make statements referring to generic “roles” in a game form Γ , to state properties of Γ that are independent of the specific agent that is assigned to that role. For instance, we previously defined $C_i(h)$ as the set of outcomes that are clinchable for an agent i at h . Below, we will sometimes write $C_r(h)$ to refer to the set of outcomes that are clinchable for the role $r \in \mathcal{R}$ at h , or $P_r(h)$ for the set of outcomes that are possible for role r . (If the role assignment function is such that $\sigma(r) = i$, then $C_i(h) = C_r(h)$, $P_i(h) = P_r(h)$, etc.) Analogously to the sets $\mathcal{A}(h)$ and $\mathcal{L}(h)$ for active agents and lurkers at a history h , we write $\mathcal{A}_R(h)$ for the set of active roles at a history h , and $\mathcal{L}_R(h)$ for the set of roles that are lurkers at h . When we want to refer to the game form with agents assigned to roles via a specific role assignment function σ_A , we will write Γ_A . In the proofs, we will often move fluidly between agents and roles; to avoid confusion, we use the notation i, j, k to refer to specific agents, and the notation r, s, t to refer to generic roles. Finally, note that while the set of lurkers at any h may differ depending on the role assignment function, the set of lurked *objects* (and the order in which they become lurked) depends only on h , and is independent of the specific agent assigned to the role that moves at h .

⁷⁴If $P_j(h) \subsetneq P_j(h')$ (because some new object became lurked between h' and h , and so disappeared as a possibility for j), then there must be a more recent subhistory $h'' \supseteq h'$ where j was re-offered the opportunity to clinch y , by definition of a millipede game (or, more primitively, by OSP).

⁷⁵For any active nonlurker i at any history h , either $P_i(h) = \mathcal{X}(h)$ or $P_i(h) = \bar{\mathcal{X}}^{\mathcal{L}}(h)$, with the former holding for at most one of the (possibly) two active non-lurkers; see Remark 7.1 of Bade and Gonczarowski (2017).

Unless otherwise specified, when we write the phrase “ i clinches x at h ” (or similar variants), what is meant is that i moves at h , takes some clinching action $a_x \in A(h)$, and receives object x at all terminal histories $\bar{h} \supseteq (h, a_x)$.

Finally, the following remark is simply a restatement of part (iii) of the definition of a lurker, but deserves special emphasis, as it will arise frequently in the arguments below.

Remark 4. If an object x has been offered to an active non-lurker at a history h (i.e., $x \in C_i^c(h)$ for some $i \in \bar{\mathcal{L}}(h)$), then $x \notin \mathcal{X}^{\mathcal{L}}((h, a))$ for any $a \in A(h)$.

Proof of Lemma A7

Let agent i^* be the first agent to clinch in game Γ_σ , which induces the ordering of the first segment of agents in step 1 of the ordering algorithm. Let $\mathcal{X}^{\mathcal{L}}(h^*) = \{x_1, \dots, x_n\}$ be the set of lurked objects at h^* (which may be empty).

Case (1): $\mathcal{A}(h) = \mathcal{L}(h) \cup \{i^*\}$.

If i^* clinches an unlurked object $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h^*)$, then, in Γ_σ , all lurkers get their lurked objects (the oldest lurker ℓ_1 gets x_1 , the second oldest lurker ℓ_2 gets x_2 , etc.), and in the resulting serial dictatorship f_σ , the agents are ordered $f_\sigma : \ell_1, \ell_2, \dots, \ell_n, i^*$. By BG Lemma A12, for each lurker ℓ_m , we have $x_m = \text{Top}(\succ_{\ell_m}, \mathcal{X} \setminus \{x_1, \dots, x_{m-1}\})$. When it is agent ℓ_m 's turn in the serial dictatorship, she is offered to choose from $\mathcal{X} \setminus \{x_1, \dots, x_{m-1}\}$, and thus selects x_m . Finally, consider agent i^* . In game Γ_σ , when she clinches y at h^* , it is unlurked. By BG Lemma A10, $\bar{\mathcal{X}}^{\mathcal{L}}(h^*) \subseteq P_i(h^*) \cup C_i^c(h^*)$, which implies that so $y = \text{Top}(\succ_{i^*}, \bar{\mathcal{X}}^{\mathcal{L}}(h^*))$. At her turn in the serial dictatorship, the set of objects remaining is precisely $\bar{\mathcal{X}}^{\mathcal{L}}(h^*)$, and so i^* will select y .

If, on the other hand, i^* clinches some lurked object x_m , then all older lurkers $\ell_1, \dots, \ell_{m-1}$ get their lurked objects in Γ_σ , and the resulting serial dictatorship begins as $f_\sigma : \ell_1, \dots, \ell_{m-1}, i^*$. By an argument equivalent to the previous paragraph, each of these agents once again gets the same object under the serial dictatorship.⁷⁶ Then, in Γ_σ , agent ℓ_m is offered to clinch anything from $\mathcal{X} \setminus \{x_1, \dots, x_m\}$. If ℓ_m takes another lurked object $x_{m'}$ for some $m' > m$, then each lurker $\ell_{m+1}, \dots, \ell_{m'-1}$ is assigned to their lurked object, and we add to the serial dictatorship order as $f_\sigma : \ell_1, \dots, \ell_{m-1}, i^*, \ell_{m+1}, \dots, \ell_{m'-1}, \ell_m$. By the same argument as above, at their turn in the resulting SD, each agent $\ell_{m+1}, \dots, \ell_{m'-1}, \ell_m$ gets the same object in the SD.⁷⁷ This process continues until someone eventually takes an unlurked object, all

⁷⁶For agent i^* , since she took a lurked object at h^* in Γ_σ , we have $x_m = \text{Top}(\succ_{i^*}, \mathcal{X})$, and thus, at her turn in the serial dictatorship, she will once again select x_m , since it is still available.

⁷⁷When it is agent ℓ_m 's turn in the SD, the set of available objects is a subset of the set of objects that were offered to her when she clinched in Γ_σ : $\mathcal{X} \setminus \{x_1, \dots, x_{m'-1}\} \subseteq \mathcal{X} \setminus \{x_1, \dots, x_m\}$. However, $x_{m'}$ belongs to both sets, and so since ℓ_m took $x_{m'}$ in Γ_σ , she will also to take it at her turn in the SD, when her offer set is smaller.

remaining lurkers are ordered, and step 1 is completed.

Case (2): $\mathcal{A}(h) = \mathcal{L}(h) \cup \{i^*, j\}$ for some $j \in \mathcal{A}(h) \setminus (\mathcal{L}(h) \cup \{i\})$.

First consider the case that i^* clinches an unlurked object $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h^*)$. If $y \notin C_j^{\subseteq}(h^*)$, then the argument is exactly the same as in Case (1) (note that j is not ordered in step 1 in this case). If $y \in C_j^{\subseteq}(h^*)$, then the step 1 partial order is $\ell_1 \tilde{\succ}^1 \dots \tilde{\succ}^1 \ell_\lambda \tilde{\succ}^1 \{i^*, j\}$. We must show that any serial dictatorship run under $f_\sigma : \ell_1, \dots, \ell_n, i^*, j, \dots$ and $f'_\sigma : \ell_1, \dots, \ell_n, j, i^*, \dots$ result in the same outcome as Γ_σ for these agents. For the lurkers, the argument is as above in either case. For i^* and j , in game Γ_σ , by construction, $y \in C_j(h')$ for some $h' \subsetneq h^*$. Let $z = \text{Top}(\succ_j, \bar{\mathcal{X}}^{\mathcal{L}}(h^*))$, and note that by Lemma A16, $z \succ_j y$. Since i clinched y at h^* , we have $y \succ_i z$. In the serial dictatorship, after all lurkers have picked, the set of remaining objects is precisely $\bar{\mathcal{X}}^{\mathcal{L}}(h^*)$. Thus, it does not matter whether i^* or j is ordered next in the serial dictatorship, as there is no conflict between them: in both cases, i^* will take y , and j will take z , and both f_σ and f'_σ give the same allocation as Γ_σ . For the case where i^* begins by clinching some lurked object $x_m \in \mathcal{X}^{\mathcal{L}}(h^*)$, we consider agent j and the lurker who, in the chain of assignments, eventually takes an unlurked object y ; otherwise, the argument is analogous.

This shows that we get the same allocation for all agents ordered in step 1 of the ordering algorithm. If all active agents at $\mathcal{A}(h^*)$ are processed in step 1 of the ordering algorithm, then we effectively have a smaller subgame on the remaining agents, and we just repeat the same argument. If not, then there is at most one active agent $j \in \mathcal{A}(h^*)$ who is not processed in step 1. Agent j has been previously offered some objects in the set $C_j^{\subseteq}(h^*)$ (note that $C_j^{\subseteq}(h^*) \subseteq \bar{\mathcal{X}}^{\mathcal{L}}(h)$). The subgame that begins after all of the agents in step 1 have clinched can equivalently be written as a Pareto efficient millipede subgame that begins with agent j being offered $C_j^{\subseteq}(h^*)$ at the “root node”. We then find the first agent to clinch (after a series of passes) in this subgame, and repeat the same argument as for step 1 above. ■

Proof of Lemma A8

We will first show the result for $k = 1$, and the proof for remaining steps will follow recursively. We first show that $h_A^1 = h_B^1$. Towards a contradiction, assume $h_A^1 \neq h_B^1$, and, wlog, $h_A^1 \subsetneq h_B^1$. Let $r = \rho(h_A^1)$ be the role associated with history h_A^1 , and define $i_A = \sigma_A(r)$ and $i_B = \sigma_B(r)$, where $i_A \neq i_B$. Since there can be at most one passing action at h_A^1 and $h_A^1 \subsetneq h_B^1$, agent i_A must clinch some x_A at h_A^1 in Γ_A , while agent i_B must pass at h_A^1 in Γ_B . Let $\mathcal{A}_R(h_A^1)$ be the set of active roles at h_A^1 . Note that by definition, $\mathcal{L}_R(h_A^1) \subseteq \mathcal{L}_R(h_B^1)$ and $\mathcal{X}^{\mathcal{L}}(h_A^1) \subseteq \mathcal{X}^{\mathcal{L}}(h_B^1)$. Also, for the constructed partial order \triangleright_A , let $g_A(i) = |j : j \triangleright_A i| + 1$.⁷⁸

⁷⁸This is almost the same as i 's picking order in the resulting serial dictatorship, except that this allows for the fact that two agents may tie under \triangleright_A , i.e., $g_A(i) = g_A(i')$ if $i \not\triangleright_A i'$ and $i' \not\triangleright_A i$

Define g_B similarly. Since we assume $\tilde{\nu}_A^1 = \tilde{\nu}_B^1$, we have $g_A(i) = g_B(i)$ for all i ordered in step 1 of the ordering algorithm applied to Γ_A and Γ_B , respectively.

Since there is a passing action at h_A^1 , we have $x_A \notin \mathcal{X}^\mathcal{L}(h_A^1)$ (by BG Lemma A11). Since i_A clinches an unlurked object $x_A \in \bar{\mathcal{X}}^\mathcal{L}(h_A^1)$ at h_A^1 , we have $x_A = \text{Top}(\succ_{i_A}, \bar{\mathcal{X}}^\mathcal{L}(h_A^1))$,⁷⁹ and also $g_A(i_A) = \lambda(h_A^1) + 1$, where $\lambda(h_A^1) = |\mathcal{L}_R(h_A^1)|$ is the number of lurkers that are present at h_A^1 . Therefore, $g_B(i_A) = \lambda(h_A^1) + 1$ as well, which implies $\mathcal{X}^\mathcal{L}(h_A^1) = \mathcal{X}^\mathcal{L}(h_B^1)$.⁸⁰ This also means that $\mathcal{L}_R(h_A^1) = \mathcal{L}_R(h_B^1)$. Let x_B be the object clinched at h_B^1 .

Case (1): $x_B \notin \mathcal{X}^\mathcal{L}(h_B^1)$.

Subcase (1).(i): $\rho(h_B^1) \neq r$. There can be at most one other active non-lurker role at h_B^1 , denoted $s \in \bar{\mathcal{L}}_R(h_B^1)$. We have $\sigma_B^{-1}(r) \neq i_A$ (or else i_A would again clinch x_A at h_A^1), and $\sigma_B^{-1}(i_A) \neq r_n$ for any lurker role $r_n \in \mathcal{L}_R(h_B^1)$ (because then $g_B(i_A) = n' < \lambda(h_A^1) + 1$, a contradiction).⁸¹ Thus, it must be that $\sigma_B(s) = i_A$, $i_{h_B^1} = i_A$, and i_A clinches x_A at h_B^1 (i.e., $x_B = x_A$). By construction of $\tilde{\nu}_B^1$, $x_A \in C_j^\subseteq(h_B^1)$, where $j = \sigma_B(r)$, and so $g_B(i_A) = g_B(j) = \lambda(h_B^1) + 1 = \lambda(h_A^1) + 1$. So, $g_A(j) = \lambda(h_A^1) + 1$; in other words, in Γ_A , when i_A clinches x_A at h_A^1 , j must be an active non-lurker at h_A^1 , and $x_A \in C_s^\subseteq(h_A^1)$, where $s = \sigma_A^{-1}(j)$ is j 's role in Γ_A . Since $\sigma_B(s) = i_A$, in game Γ_B , there is some $h' \subsetneq h_A^1$ such that $x_A \in C_{i_A}(h')$ and i passes at h' . Let $\bar{x} = \text{Top}(\succ_{i_A}, \bar{\mathcal{X}}^\mathcal{L}(h_A^1))$, and note that by Lemma A16, $\bar{x} \succ_{i_A} x_A$. However, we saw above that $x_A = \text{Top}(\succ_{i_A}, \bar{\mathcal{X}}^\mathcal{L}(h_A^1))$, which is a contradiction.

Subcase (1).(ii): $\rho(h_B^1) = r$. In game Γ_B , $i_{h_B^1} = i_B$, and, at h_B^1 , i_B clinches some $x_B \neq x_A$. Since x_B is unlurked at h_B^1 , we have $x_B = \text{Top}(\succ_{i_B}, \bar{\mathcal{X}}^\mathcal{L}(h_B^1))$. Since $g_B(i_A) = \lambda(h_A^1) + 1$ (as required by $\tilde{\nu}_A^1 = \tilde{\nu}_B^1$), we have $i_A \in \mathcal{A}(h_B^1)$ and $x_B \in C_{i_A}^\subseteq(h_B^1)$ in Γ_B . This again implies $g_A(i_A) = g_A(i_B) = g_B(i_A) = g_B(i_B) = \lambda(h_A^1) + 1$, i.e., i_A and i_B tie in both $\tilde{\nu}_A^1$ and $\tilde{\nu}_B^1$. Let $s = \sigma_B^{-1}(i_A)$. Since $g_A(i_B) = \lambda(h_A^1) + 1$, i_B must have been an active nonlurker at h_A^1 in Γ_A , which means that $\sigma_A(i_B) = s$. Therefore, in game Γ_A , i_B passes at some history $h' \subsetneq h_A^1$ such that $x_B \in C_s(h')$. An argument equivalent to the previous paragraph applied to i_B again reaches a contradiction.

Case (2): $x_B \in \mathcal{X}^\mathcal{L}(h_B^1)$. Note that in this case, since a lurked object is clinchable at h_B^1 , there is no passing action at h_B^1 , by BG Lemma A11. Further, the role/agent who moves at h_B^1 satisfies the conditions of the terminator t defined in BG Lemma A14; denote $\rho(h_B^1) = t$, and note that $C_t^\subseteq(h_B^1) = \mathcal{X}(h_B^1) = \mathcal{X}$. Also, recall from the discussion before Case (1) that $g_A(i_A) = \lambda(h_A^1) + 1$, where $\lambda(h_A^1) = |\mathcal{L}_R(h_A^1)|$, and therefore, $g_B(i_A) = \lambda(h_A^1) + 1$ as well.

⁷⁹This follows because $\bar{\mathcal{X}}^\mathcal{L}(h_A^1) \subseteq P_{i_A}(h_A^1) \cup C_{i_A}^\subseteq(h_A^1)$, by Lemma A10.

⁸⁰To see this, note that if $\mathcal{X}^\mathcal{L}(h_B^1) \supsetneq \mathcal{X}^\mathcal{L}(h_A^1)$, then $x_{\lambda(h_A^1)+1} = x_A$, i.e., the $(\lambda(h_A^1) + 1)^{\text{th}}$ lurked object must be x_A (because the ordering algorithm puts the agent who receives $x_{\lambda(h_A^1)+1}$ as the $(\lambda(h_A^1) + 1)^{\text{th}}$ agent in the ordering, and we know $g_B(i_A) = \lambda(h_A^1) + 1$). Thus, let $h' \supsetneq h_A^1$ be the history where x_A first becomes lurked. Note that $x_A \in C_r^\subseteq(h')$. However, role r is still a non-lurker at h' , and so x_A cannot become lurked (see the definition of lurker /Remark 4).

⁸¹Note that in this case, x_B is unlurked, and so all lurkers are immediately assigned to their lurked objects.

Subcase (2).(i): $t = r$. In this case, in Γ_A , when i_A clinches x_A at h_A^1 , we have $x_A = \text{Top}(\succ_i, \mathcal{X})$ (because $\sigma_A^{-1}(i_A) = t$ and i_A chose to clinch first). Now, since $\sigma_B^{-1}(i_A) \neq t$, the only way for $g_B(i_A) = \lambda(h_A^1) + 1$ is for i_A to be the active non-lurker at h_B^1 that does not clinch, and $y \in C_{i_A}^{\subset}(h_B^1)$, where y is the unlurked object chosen by some lurker $i_\ell \in \mathcal{L}(h_B^1)$ in the assignment chain initiated when i_B selected x_B at h_B^1 . Let $s = \sigma_B^{-1}(i_A)$. Since $i_\ell \in \mathcal{L}(h_B^1)$, and chose y at her turn, we have $y = \text{Top}(\succ_{i_\ell}, \bar{\mathcal{X}}^{\mathcal{L}}(h_B^1))$. Note that $g_B(i_\ell) = \lambda(h_A^1) + 1$, and so, since $\tilde{\delta}_A^1 = \tilde{\delta}_B^1$, we have $g_A(i_\ell) = \lambda(h_A^1) + 1$ as well. This is only possible if $\sigma_A^{-1}(i_\ell) = s$, and $x_A \in C_s^{\subset}(h_A^1)$. But then, in game Γ_B , agent $i_A (= \sigma_B(s))$ was offered to clinch x_A at some history $h' \subsetneq h_B^1$. Since $x_A = \text{Top}(\succ_i, \mathcal{X})$, i_A clinches at h' in Γ_B , which is a contradiction.

Subcase (2).(ii): $t \neq r$. In this case, in Γ_A , when i_A clinches x_A at h_A^1 , since $\sigma_A(t) \neq i_A$, we have $x_A \notin C_t^{\subset}(h_A^1)$, by BG Lemma A14. Therefore, $g_A(i_A) = \lambda(h_A^1) + 1$, and $g_A(i') \neq \lambda(h_A^1) + 1$ for all other $i' \neq i_A$ ordered in step 1 (in other words, i_A does not tie with another agent in $\tilde{\delta}_A^1$), and so the same is true for g_B .⁸² Since the first agent to clinch in Γ_B is $\sigma_B(t) = j \neq i_A$ who clinches some lurked object $x_B \in \mathcal{X}^{\mathcal{L}}(h_B^1)$ and $\sigma_B(r) = i_B \neq i_A$, $\sigma_B^{-1}(i_A) = r_n$ for some lurker role r_n that lurks object $x_{n'} \in \mathcal{X}^{\mathcal{L}}(h_B^1)$. Now, when i_A eventually clinches x_A (after someone else has selected $x_{n'}$ in the chain of lurker assignments), $x_A \in C_{i_B}^{\subset}(h_B^1)$, where $i_B = \sigma_B(r)$ (since $h_B^1 \supseteq h_A^1$ and $x_A \in C_r(h_A^1)$), which implies that $g_B(i_B) = \lambda(h_A^1) + 1$, i.e., i_B ties with i_A —a contradiction.

Thus far, we have shown that $\tilde{\delta}_A^1 = \tilde{\delta}_B^1$ implies $h_A^1 = h_B^1$, or, in other words, if step 1 of the ordering algorithm produces the same ordering, then step 1 must be initiated at the same history in Γ_A and Γ_B . Next, we show that $\sigma_A(r') = \sigma_B(r')$ for all r' that are ordered in step 1 of Γ_A and Γ_B .

Define $h^1 := h_A^1 = h_B^1$. Let $\mathcal{L}_R(h^1) = \{r_1, \dots, r_{\lambda(h^1)}\}$ be the set of lurker-roles at h^1 , and $\mathcal{X}^{\mathcal{L}}(h^1) = \{x_1, \dots, x_{\lambda(h^1)}\}$ the set of lurked objects. Notice that, since $h_A^1 = h_B^1$, the lurked objects and active lurker-roles are equivalent in both Γ_A and Γ_B . Towards a contradiction, assume that $\sigma_A(r') \neq \sigma_B(r')$ for some r' that is ordered in step 1. Letting $r_0 = \rho(h^1)$, write

$$\sigma_A(r_0) \rightarrow x_{a_1} \rightarrow \sigma_A(r_{a_1}) \rightarrow x_{a_2} \rightarrow \dots \rightarrow \sigma_A(r_{a_M}) \rightarrow y_A \quad (1)$$

to represent the chain of clinching that is initiated in Γ_A by agent $\sigma_A(r_0)$ at h^1 : agent $\sigma_A(r_0)$ clinches some (possibly lurked) object x_{a_1} , the agent $\sigma_A(r_{a_1})$ who was lurking x_{a_1} clinches lurked object x_{a_2} , etc., until eventually agent $\sigma_A(r_{a_M})$ ends the chain by clinching some

⁸²Note that since $\mathcal{L}_R(h_A^1) = \mathcal{L}_R(h_B^1)$, there cannot be any additional role $r' \notin \mathcal{L}_R(h_A^1) \cup \{r, t\}$ that is active at h_A^1 .

unlurked object y_A .⁸³ Similarly, for Γ_B , write

$$\sigma_B(r_0) \rightarrow x_{b_1} \rightarrow \sigma_B(r_{b_1}) \rightarrow x_{b_2} \rightarrow \cdots \rightarrow \sigma_B(r_{b_M}) \rightarrow y_B. \quad (2)$$

We will show that chains (1) and (2) above are in fact equivalent: $\sigma_A(r_0) = \sigma_B(r_0)$ and $\sigma_A(r_{a_m}) = \sigma_B(r_{b_m})$ for all m . Since any lurked object $x_n \in \mathcal{X}^{\mathcal{L}}(h^1)$ that does not appear in the above chain must be assigned to its lurker $\sigma_A(r_n)$, this will deliver the result.

First, note that if $x_{a_1} = x_{b_1}$ and $x_{a_1} \notin \mathcal{X}^{\mathcal{L}}(h^1)$, then both (1) and (2) begin with the same agent taking the same (unlurked) object. Therefore, all lurkers are immediately assigned to their lurked objects. If there is another active non-lurker role $s \in \bar{\mathcal{L}}_R(h^1)$, then the agent in role s clinches his object favorite remaining object. In either case, it is clear that $\sigma_A(r') = \sigma_B(r')$ for all roles r' ordered in step 1 of the ordering algorithm. Thus, assume that $x_{a_1} \neq x_{b_1}$, and therefore, $\sigma_A(r_0) \neq \sigma_B(r_0)$.

Claim A2. At least one of x_{a_1} or x_{b_1} is lurked at h^1 ; i.e., $\mathcal{X}^{\mathcal{L}}(h^1) \cap \{x_{a_1}, x_{b_1}\} \neq \emptyset$.

Proof of claim. Assume that $x_{a_1}, x_{b_1} \notin \mathcal{X}^{\mathcal{L}}(h^1)$. For shorthand, define $\sigma_A(r_0) = i_A$, and $\sigma_B(r_0) = i_B$. Then, since $\tilde{\nu}_A^1 = \tilde{\nu}_B^1$, we have $g_A(i_A) = g_A(i_B) = g_B(i_A) = g_B(i_B) = \lambda(h^1) + 1$, i.e., i_A and i_B tie under both $\tilde{\nu}_A^1$ and $\tilde{\nu}_B^1$. This means that there is another active non-lurker role $s \in \mathcal{A}_R(h^1) \setminus (\mathcal{L}_R(h^1) \cup \{r_0\})$, and $\sigma_A(r_0) = \sigma_B(s) = i_A$, $\sigma_A(s) = \sigma_B(r_0) = i_B$. Further, $x_{a_1}, x_{b_1} \in C_s^{\mathcal{C}}(h^1)$.

If $\mathcal{L}(h^1) = \emptyset$, then $\mathcal{X} = P_{r_0}(h^1) \cup C_{r_0}^{\mathcal{C}}(h^1)$, which implies that $x_{a_1} = \text{Top}(\succ_{i_A}, \mathcal{X})$ and $x_{b_1} = \text{Top}(\succ_{i_B}, \mathcal{X})$. But, this means that in Γ_A , i_B will clinch x_{b_1} at some $h' \subsetneq h^1$ (in particular, the earliest h' such that $x_{b_1} \in C_s(h')$), a contradiction. Therefore, $\mathcal{L}(h^1) \neq \emptyset$. Now, by BG Lemma A10, $x_{a_1} = \text{Top}(\succ_{i_A}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1))$ and $x_{b_1} = \text{Top}(\succ_{i_B}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1))$. Since $x_{b_1} \in C_s^{\mathcal{C}}(h^1)$, there is some history $h' \subsetneq h^1$ such that in game Γ_A , $x_{b_1} \in C_{i_B}(h')$, and i_B passes at h' . By Lemma A16, $\text{Top}(\succ_{i_B}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1)) \succ_{i_B} x_{b_1}$, which is a contradiction. ■

By the previous claim, at least one (possibly both) of x_{a_1} and x_{b_1} are lurked objects; wlog, assume that $x_{a_1} \in \mathcal{X}^{\mathcal{L}}(h^1)$. Since $C_{r_0}(h^1)$ contains a lurked object, BG Lemma A14 implies $r_0 = t$ is the terminator role. Consider agent $\sigma_A(r_{a_M}) = i'$, and note that $x_{a_M} \succ_{i'} y_A = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_M}\})$ where x_{a_M} is the object i' lurks at h^1 in Γ_A .

We first claim that $\sigma_B^{-1}(i') = r_{a_M}$ and $y_A = y_B$. To see this, first note that $\sigma_B^{-1}(i') \neq r_{n''}$ for any $n'' > a_M$. Indeed, if this were the case, this would imply that $x_{n''} = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{n''-1}\})$, where $x_{n''}$ is the object lurked by role $r_{n''}$. However, this contradicts $y_A = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_M}\})$. Next, we show that $\sigma_B^{-1}(i') \neq r_{n''}$ for any $n'' < a_M$ either. In game Γ_B , when i' becomes a lurker for $x_{n''}$ at some h' , he eventually must get no worse

⁸³Recall that any lurked object that does not appear in the chain is assigned to its lurker. For example, if $a_1 < a_2$, then $x_{a'}$ is assigned to $\sigma_A(r_{a'})$ for all a' such that $a_1 < a' < a_2$.

than his second-best choice from the set $P_{i'}(h') = \mathcal{X} \setminus \{x_1, \dots, x_{n''-1}\}$. Since $x_{a_M} \in P_{i'}(h')$, we have $x_{n''} \succ_{i'} x_{a_M} \succ_{i'} y_A$, and i' can do no worse than x_{a_M} , which means he cannot end up with y_A —a contradiction. The final case to consider is $\sigma_B^{-1}(i') = r'$ for some $r' \in \bar{\mathcal{L}}_R(h^1)$. We cannot have $r' = r_0$ (since r_0 is the terminator role, i' would then be able to clinch her top choice at h^1 in Γ_B , and $Top(\succ_{i'}, \mathcal{X}) \neq y_A$). Thus, r' must be the (unique) other non-lurker role that is active at h^1 : $r' = \mathcal{A}_R(h^1) \setminus (\mathcal{L}_R(h^1) \cup \{r_0\})$. Recall that, for this role, $P_{r'}(h^1) = \bar{\mathcal{X}}^{\mathcal{L}}(h^1)$. Further, $y_A \notin C_{r'}^{\subseteq}(h^1)$ (or else i' would have clinched y_A at some strict subhistory of h^1), and $y_B \in C_{r'}^{\subseteq}(h^1)$ (or else i' would not be ordered in step 1 of the ordering algorithm under Γ_B). However, the former implies that $|j : g_A(j) = \lambda(h^1) + 1| = 1$, while the latter implies $|j : g_B(j) = \lambda(h^1) + 1| = 2$, which contradicts $\tilde{\delta}_A^1 = \tilde{\delta}_B^1$. Therefore, $\sigma_A^{-1}(i') = \sigma_B^{-1}(i') = r_{a_M}$. Finally, note that $\sigma_A(r_{a_M}) = \sigma_B(r_{a_M})$ further implies $y_A = y_B$ and $r_{a_M} = r_{b_M}$. Indeed, if not, then the final person to clinch in chain 2 is some $\sigma_B(r_{b_M}) = j \neq i'$. However, agent i' is a lurker in Γ_B for x_{a_M} that was not previously taken by any other agent in chain 2, and thus, i' is assigned to x_{a_M} , which is a contradiction.

Next, consider agent $\sigma_A(r_{a_{M-1}}) = i'$ in chain 1, i.e., i' lurks $x_{a_{M-1}}$ in Γ_A and eventually ends up with (lurked) object x_{a_M} . By construction of the chain, $x_{a_M} = Top(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_{M-1}}\})$. Similar to the previous paragraph, $\sigma_B^{-1}(i') \neq r_{n''}$ for any $n'' > a_M - 1$. Indeed, if this were true, then $x_{n''} = Top(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{n''-1}\})$. If $n'' < a_M$, $x_{n''} \succ_{i'} x_{a_M}$, which contradicts $x_{a_M} = Top(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_{M-1}}\})$. If $n'' > a_M$, then x_{a_M} is not possible for i' (BG Lemma A12), which is also a contradiction. Finally, $n'' \neq a_M$, since we already have shown that $\sigma_B(r_{a_M}) = \sigma_A(r_{a_M})$. Similarly, $\sigma_B^{-1}(i') \neq r_{n''}$ for any $n'' < a_M - 1$, either, since this would imply that $x_{n''} \succ_{i'} x_{a_{M-1}} \succ_{i'} x_{a_M}$, and $x_{a_{M-1}} \in P_{i'}(h')$ at the history h' where i' became a lurker for $x_{n''}$. Since i' cannot do any worse than his second-best choice from $P_{i'}(h')$, we have a contradiction. The last case to consider is $\sigma_B^{-1}(i') = r'$ for $r' = \mathcal{A}_R(h^1) \setminus (\mathcal{L}_R(h^1) \cup \{r_0\})$.⁸⁴ But, for this role, $P_{r'}(h^1) = \bar{\mathcal{X}}^{\mathcal{L}}(h^1)$, and thus, no lurked objects are possible for i' , which is a contradiction. Therefore, $\sigma_A(r_{a_{M-1}}) = \sigma_B(r_{a_{M-1}})$. As in the previous paragraph, this also implies that $r_{a_{M-1}} = r_{b_{M-1}}$ and $x_{a_M} = x_{b_M}$.

The same argument can be repeated to show that $x_{a_m} = x_{b_m}$ and $\sigma_A(r_{a_m}) = \sigma_B(r_{b_m})$ for all $m = 1, \dots, M$. The final case to consider is role r_0 . Let $\sigma_A(r_0) = i'$. Since i' starts the chain of assignments at h^1 by taking some lurked object $x_{a_1} \in \mathcal{X}^{\mathcal{L}}(h^1)$, we have $x_{a_1} = Top(\succ_{i'}, \mathcal{X})$. Once again, we cannot have $\sigma_B^{-1}(i') = r_{n''}$ for any $n'' < a_1$, as this would imply that $x_{n''} \succ_{i'} x_{a_1}$, which is a contradiction. We also cannot have $\sigma_B^{-1}(i') = r_{a_1}$, since we have already shown that $\sigma_B(r_{a_1}) = \sigma_A(r_{a_1})$. Further, we cannot have $\sigma_B^{-1}(i') = r_{n''}$ for any $n'' > a_1$, since x_{a_1} would not be possible for i' . Last, we cannot have $\sigma_B^{-1}(i') = r'$ for $r' = \mathcal{A}_R(h^1) \setminus (\mathcal{L}_R(h^1) \cup \{r_0\})$, since no lurked objects are possible for the agent in role r' .

⁸⁴The case where $\sigma_B^{-1}(i') = r_0$ can be dispensed with similarly as in the previous paragraph.

Therefore, $\sigma_B^{-1}(i') = r_0$, and chains 1 and 2 are equivalent.

To summarize: We have shown that if we have two role assignment functions σ_A, σ_B such that $\tilde{\sigma}_A^1 = \tilde{\sigma}_B^1$, then (i) $\sigma_A^{-1}(i) = \sigma_B^{-1}(i)$ for all agents i that are ordered in step 1 of the ordering algorithm and (ii) at the conclusion of the chain of clinching initiated by the first agent to start the chain at h_A^1/h_B^1 , we will end up at the same history in Γ_A as in Γ_B .⁸⁵ At this point, we have a smaller subgame consisting of the agents and objects that were unmatched in the first round. This subgame is another Pareto efficient millipede game (that may begin with the unique unmatched agent from step 1 carrying over his endowment $C_j^{\subseteq}(h^1)$, if such an agent exists), and so we simply repeat the above arguments for each round $k = 1, \dots, K$. ■

Proof of Lemma A9

Assume there are three permutations σ_A, σ_B , and σ_C that deliver (initial) partial orderings $\triangleright_A, \triangleright_B, \triangleright_C$ as in the statement. We'll show that these 3 conditions lead to a contradiction.

As with Lemma A8, we show this first for the case that under \triangleright_A , all agents $\{i_1, \dots, i_n, i, j\}$ are processed in step 1 of the ordering algorithm, and the argument for later steps will be equivalent. Let Γ_A, Γ_B , and Γ_C denote the specific games under role assignments σ_A, σ_B , and σ_C respectively. Further, let h_A^*, h_B^* , and h_C^* be the first history at which an object is clinched in the respective games, following a sequence of passes.⁸⁶ In particular, this means that in \triangleright_A , agents $\{i_1, \dots, i_n\}$ are getting lurked objects $\mathcal{X}^{\mathcal{L}}(h_A^*) = \{x_1, \dots, x_n\}$, while agents i and j are getting some unlurked objects, $y, z \notin \mathcal{X}^{\mathcal{L}}(h_A^*)$, respectively. By construction, one of i or j must be an active non-lurker at h_A^* who is not called to move at h_A^* ; without loss of generality, assume that this is j . For notational purposes, denote by $i_{h_A^*} \in \{i_1, \dots, i_n, i\}$ the agent who moves (and clinches) at h_A^* , thereby starting the chain of lurker assignments that ends with i clinching y followed by j clinching z (note that $i_{h_A^*} \notin \mathcal{L}(h_A^*)$, by Lemma A15, and it is possible that $i_{h_A^*} = i$). The structure of \triangleright_A implies that $y \in C_j^{\subseteq}(h_A^*)$.

Case 1: $\mathcal{L}(h_A^*) = \emptyset$

In this case, by definition of \triangleright_A , the set of active agents at h_A^* is $\mathcal{A}(h_A^*) = \{i, j\}$, where $i_{h_A^*} = i$. For notational purposes, let $s = \sigma_A^{-1}(i)$ and $s' = \sigma_A^{-1}(j)$ be the roles assigned to agents i and j in game Γ_A , and note that $y \in C_{s'}^{\subseteq}(h_A^*)$. Further, both i and j are getting their first choice objects, i.e., $Top(\succ_i, \mathcal{X}) = y$ and $Top(\succ_j, \mathcal{X}) = z$.

⁸⁵This follows because the chain of clinching starts at the same history in both games, and all agents who are ordered in step 1 are in the same roles, so will take the same actions. While there may be some other role r' that is not ordered in step 1 and $\sigma_A(r') \neq \sigma_B(r')$, this agent must have passed every time she was called to move at a history $h' \subsetneq h^1$, and is not called to move in the chain of lurker assignments, and so at the end of the chain, we will still end up at the same history to begin the next round/subgame.

⁸⁶That is, $h_\alpha^* = (h_\emptyset, a^*, \dots, a^*)$ for $\alpha = A, B, C$, though the number of passes (a^* 's) may vary.

Now, consider σ_B , where the ordering algorithm produces $i \triangleright_B j$. By Remark 4, y cannot be the first lurked object in the game, and thus, for i to be ordered first without ties according to \triangleright_B , we must have $\mathcal{X}^{\mathcal{L}}(h_B^*) = \emptyset$ and $i_{h_B^*} = i$. These facts imply that $\sigma_B^{-1}(i) = s'$.

Now, consider σ_C , which begins $j \triangleright_C i \cdots$. There are two subcases, depending on whether j is a lurker at h_C^* or not.

Subcase 1.(i): $j \in \mathcal{L}(h_C^*)$. In this case, by construction, j is the first lurker of the game and z is the first lurked object. Further, $h_C^* \supseteq h_A^*$. Now, in order for i to be the (unique) next agent added to \triangleright_C , either (i) $y \in \mathcal{X}^{\mathcal{L}}(h_C^*)$ and, in particular, y is the second object to become lurked in the game or (ii) $\mathcal{X}^{\mathcal{L}}(h_C^*) = \{z\}$ and i clinches y at h_C^* . But, by Remark 4, y cannot be the second lurked object of the game, because y was previously offered to both roles s and s' , and, even after z becomes lurked by j at some history h' , we will still have $y \in C_{r'}^{\subseteq}(h')$ for the role $r' \in \{s, s'\}$ such that $\sigma_C^{-1}(j) \neq r'$. Therefore, i must clinch y at $h_C^* \supsetneq h_A^*$. Now, we have $\sigma_C^{-1}(i) \neq s, s'$ (because $Top(\succ_i, \mathcal{X}) = y$, and so i would have clinched y earlier along the path to h_C^* , since it has been previously offered to both roles). Therefore, agent $k = \sigma_C(r')$ is an active non-lurker at h_C^* such that $y \in C_k^{\subseteq}(h_C^*)$, and so, by construction of \triangleright_C , we have $j \triangleright_C \{i, k\} \triangleright_C \cdots$, which is a contradiction.

Subcase (ii): $j \notin \mathcal{L}(h_C^*)$. In this case, $\mathcal{L}(h_C^*) = \emptyset$ and $i_{h_C^*} = j$ (since j is ordered first without ties). Further, $\sigma_C^{-1}(j) \in \{s, s'\}$. If $\sigma_C^{-1}(j) = s'$, then $\sigma_C(s) = k \neq i$ (or else we are back to σ_A) and $h_C^* \supset h_A^*$. Thus, $y \in C_k^{\subseteq}(h_A^*)$, which implies that i cannot be the next agent added to \triangleright_C uniquely—a contradiction.⁸⁷ The last case is $\sigma_C^{-1}(j) = s$. Note that $z \notin C_s^{\subseteq}(h_B^*)$,⁸⁸ and so $h_C^* \supset h_B^*$. This further implies that $y \in C_{s'}^{\subseteq}(h_C^*)$ and so $\sigma_C(s') = k \neq i$ (since otherwise, i would have clinched y prior to h_C^* , because $Top(\succ_i, \mathcal{X}) = y$). By an argument similar to footnote 87, i cannot be the next agent added to the order uniquely, which is again a contradiction.

This completes the argument for Case 1.

Case 2: $\mathcal{L}(h_A^*) \neq \emptyset$.

Now, we consider the case where there are lurkers at h_A^* (and hence also lurked objects, $\mathcal{X}^{\mathcal{L}}(h_A^*) \neq \emptyset$). By definition of \triangleright_A , we have $\mathcal{A}_R(h_A^*) = \{r_1, \dots, r_n, s, s'\}$, where $r_1, \dots, r_n \in \mathcal{L}_R(h_A^*)$ are lurker roles, and $s, s' \in \bar{\mathcal{L}}_R(h_A^*)$ are non-lurker roles. Let $\rho(h_A^*) = s$ be the non-lurker role that moves at h_A^* (and thus clinches in Γ_A).

⁸⁷For agent i to be ordered next in \triangleright_C without ties, she must be the first agent ordered in step 2 of the ordering algorithm. By Remark 4 again, y cannot be the next lurked object in the game, which means that i must clinch y at some $h' \supsetneq h_C^*$ such that $\mathcal{L}(h') = \emptyset$. But, $y \in C_k^{\subseteq}(h')$ and so, by construction of \triangleright_C , she will tie with agent k .

⁸⁸If $z \in C_s^{\subseteq}(h_B^*)$, then, $\sigma_B^{-1}(s) = k \neq j$ (if $\sigma_B^{-1}(s) = j$, then j would clinch z prior to h_B^* , since $z = Top(\succ_j, \mathcal{X})$). But, by the same argument in footnote 87, j could not be the next agent added to \triangleright_B uniquely, a contradiction.

Claim A3. At h_A^* , y has been previously offered to both active non-lurker roles: $y \in C_s^{\subseteq}(h_A^*)$ and $y \in C_{s'}^{\subseteq}(h_A^*)$.

Proof of claim. If $\sigma_A(s) = i$, then it is obvious that $y \in C_s^{\subseteq}(h_A^*)$ by definition. If $\sigma_A(s) = i_{n'}$, then at h_A^* , agent $i_{n'}$ clinches a lurked object $x_{n'} \in C_s(h_A^*)$, which initiates the chain of lurker assignments.⁸⁹ This implies that there is no passing action at h_A^* (by Lemma A11) and that s is the terminator role t defined in BG Lemma A14. Therefore, by BG Lemma A14, $y \in C_s^{\subseteq}(h_A^*)$. That $y \in C_{s'}^{\subseteq}(h_A^*)$ for the other non-lurker role s' follows immediately from the construction of \triangleright_A . ■

Claim A4. At h_B^* , $\mathcal{X}^{\mathcal{L}}(h_A^*) \subseteq \mathcal{X}^{\mathcal{L}}(h_B^*)$. Similarly, at h_C^* , $\mathcal{X}^{\mathcal{L}}(h_A^*) \subseteq \mathcal{X}^{\mathcal{L}}(h_C^*)$.

Proof of claim. Assume not, and let $\mathcal{X}^{\mathcal{L}}(h_B^*) = \{x_1, \dots, x_{n'}\} \subsetneq \mathcal{X}^{\mathcal{L}}(h_A^*) = \mathcal{X}^{\mathcal{L}}(h_B^*) \cup \{x_{n'+1}, \dots, x_n\}$. (Recall that following a sequence of passes to start the game, there is a unique order in which objects will become lurked that is independent of the role assignment function. Since h_A^* and h_B^* are by definition the first histories at which an object is clinched in their respective games, at least one of $\mathcal{X}^{\mathcal{L}}(h_A^*) \subseteq \mathcal{X}^{\mathcal{L}}(h_B^*)$ or $\mathcal{X}^{\mathcal{L}}(h_B^*) \subseteq \mathcal{X}^{\mathcal{L}}(h_A^*)$ must hold). Agent $i_{h_B^*}$ (the agent who clinches at h_B^*) cannot clinch $x_{\tilde{n}}$ for any $\tilde{n} \leq n'$. To see why, note that this would imply that agent $i_{h_B^*}$ is offered a previously lurked object at h_B^* . By BG Lemma A11, there is no passing action at h_B^* , which contradicts $x_{n'+1} \in \mathcal{X}^{\mathcal{L}}(h_A^*)$. Thus, the only other possibility consistent with \triangleright_B is that $i_{h_B^*} = i_{n'+1}$, who clinches $x_{n'+1}$ (which is unlurked) at h_B^* , i.e., $x_{n'+1} \in C_{i_{n'+1}}(h_B^*)$. But, $x_{n'+1}$ is the (unique) next object to become lurked in game form Γ following a sequence of passes from h_B^* , which implies that $x_{n'+1} \notin C_{i_{n'+1}}(h_B^*)$, by Remark 4. An analogous argument applies for h_C^* . ■

Claim A5. $\mathcal{X}^{\mathcal{L}}(h_A^*) = \mathcal{X}^{\mathcal{L}}(h_B^*) = \mathcal{X}^{\mathcal{L}}(h_C^*)$.

Proof of claim. Given Claim A4, it is sufficient to show $\mathcal{X}^{\mathcal{L}}(h_B^*), \mathcal{X}^{\mathcal{L}}(h_C^*) \subseteq \mathcal{X}^{\mathcal{L}}(h_A^*)$. Let $\mathcal{X}^{\mathcal{L}}(h_A^*) = \{x_1, \dots, x_n\}$, and note again that the order in which objects become lurked following a series of passes to start the game is unique and independent of the role assignment function. Thus, it is sufficient to consider the next object that can become lurked, x_{n+1} , and show that $x_{n+1} \notin \mathcal{X}^{\mathcal{L}}(h_B^*)$ (resp. $x_{n+1} \notin \mathcal{X}^{\mathcal{L}}(h_C^*)$). Thus, assume that $x_{n+1} \in \mathcal{X}^{\mathcal{L}}(h_B^*)$, and note that this implies $h_B^* \supseteq h_A^*$. By construction of \triangleright_B , we must have $x_{n+1} = y$ (the object received by i); indeed, if $x_{n+1} \neq y$, then, the agent, say k , who receives x_{n+1} will be such that $k \triangleright_B i$, which is a contradiction. However, since $h_B^* \supseteq h_A^*$, y has previously been offered to both active non-lurkers at h_B^* (from the construction of \triangleright_A). Thus, by Remark 4, y cannot become the next object lurked, i.e., $x_{n+1} \neq y$ —a contradiction.

Next consider Γ_C , and assume that $x_{n+1} \in \mathcal{X}^{\mathcal{L}}(h_C^*)$. Let r_{n+1} be the role that lurks x_{n+1} , and h_{n+1} be the history at which role r becomes a lurker for x_{n+1} (i.e., role r_{n+1}

⁸⁹By definition, $\sigma_A(s) \neq j$, and so this exhausts all possibilities.

passes at h_{n+1} , and becomes a lurker at $h' = (h_{n+1}, a^*)$. Note that $h' \supset h_A^*$. Further, from what we know about the structure of the game tree Γ from σ_A , there is another active non-lurker role at h_{n+1} , denoted \tilde{r} , and we have $y \in C_{r_{n+1}}^{\subseteq}(h_{n+1})$ and $y \in C_{\tilde{r}}^{\subseteq}(h_{n+1})$. Now, since $i_1 \triangleright_C \cdots \triangleright_C i_n \triangleright_C j \triangleright_C i \cdots$, it must be that $x_{n+1} = z$ (the object received by j). Since i is uniquely ordered immediately after j , we have $i_{n+1} = j$, and $i_{h_C^*} = i$ who clinches y at h_C^* . (The only other possibility is that there is another lurked object at h_C^* , $x_{n+2} = y$, but, by Remark 4, this is impossible, since y has been previously offered to both active nonlurkers at h_{n+1}). Since i is ordered uniquely, $\sigma_C(\tilde{r}) = i$, by an argument equivalent to footnote 87. Now, this implies that i was previously offered y at some $h'' \subsetneq h_C^*$, and chose to pass, which implies $Top(\succ_i, P_i(h'')) = \bar{x} \succ_i y$. Letting h'' be the most recent subhistory such that $y \in C_i(h'')$ and i passes, and noting that i chose to clinch y at h_C^* , we conclude that $\bar{x} \notin P_i(h_C^*)$ and $\bar{x} = z$.⁹⁰ But, $z \notin \mathcal{X}^{\mathcal{L}}(h_A^*)$, which contradicts that she clinches y in Γ_A .⁹¹ ■

Claim A5 thus implies that $\mathcal{L}_R(h_A^*) = \mathcal{L}_R(h_B^*) = \mathcal{L}_R(h_C^*)$. Further, we know from BG Lemma A13 that there can be at most two active non-lurker roles at any history, which we will denote s and s' . By construction, we know that both of these roles are in fact active at h_A^* , i.e., $\mathcal{A}_R(h_A^*) = \mathcal{L}_R(h_A^*) \cup \{s, s'\}$. For h_B^* , it is possible that only one of s or s' are active (but this can only occur if $h_B^* \subsetneq h_A^*$). A similar remark applies to h_C^* .

Claim A6. $h_B^*, h_C^* \subsetneq h_A^*$.

Proof of claim. First, assume $h_A^* \subseteq h_B^*$. Then, when agent i clinches y in Γ_B (either at h_B^* , or in the chain of lurker assignments that follows), it has already been offered to both of the agents in roles s and s' , including the (unique) active non-lurker that does not move at $i_{h_B^*}$, say agent k , and so k is ordered in step 1 (and in particular, k will “tie” with i), which is a contradiction to the definition of \triangleright_B .

Next, assume that $h_A^* \subseteq h_C^*$. The agents processed in step 1 of the ordering algorithm applied to Γ_C are $\{i_1, \dots, i_n, j\}$ (a set that does not include i), and the chain ends when j clinches z . Since $h_C^* \supseteq h_A^*$, both $s, s' \in \mathcal{A}_R(h_C^*)$, and $y \in C_s^{\subseteq}(h_C^*)$ and $y \in C_{s'}^{\subseteq}(h_C^*)$. Since i is the next agent ordered in \triangleright_C without ties (in step 2 of the ordering algorithm), we have $\sigma_C(s') = i$ and $\sigma_C(s) = j$ (by an argument equivalent to footnote 87). Since i clinches $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*)$ in Γ_A , we have $y = Top(\succ_i, \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*))$. Since $\sigma_C(s') = i$, there is some $h' \subsetneq h_A^*$ such that $y \in C_i(h')$ and i passes at h' in Γ_C . By Lemma A16, $Top(\succ_i, \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*)) = \bar{x} \succ_i y$,

⁹⁰Each time a new object becomes impossible for i (due to becoming lurked by another agent), i must once again be offered the opportunity to clinch y , by definition of a millipede game. Agent i must have passed at all such opportunities (including h'') up to h_C^* , which implies that $\bar{x} = z$.

⁹¹There are actually two subcases here: if $i_{h_A^*} \neq i$, then i must be a lurker for some $x_{n'}$. At some point in the lurker assignment chain, someone (either $i_{h_A^*}$, or an earlier lurker) takes $x_{n'}$; since z is still unlurked at that point, it is possible for i . Similarly, if $i_{h_A^*} = i$, then z is still unlurked at h_A^* , which again contradicts that i clinches y .

which is a contradiction. ■

Claim A7. At h_B^* , we have $\mathcal{X}^{\mathcal{L}}(h_B^*) \cap C_{r'}(h_B^*) = \emptyset$, where $\rho(h_B^*) = r'$. The same holds at h_C^* .

Proof of claim. By Claim A6, $h_B^* \subsetneq h_A^*$. This implies that there must be a passing action at h_B^* , which means that $x_{n'} \notin C_{r'}(h_B^*)$ for any $x_{n'} \in \mathcal{X}^{\mathcal{L}}(h_B^*)$ by BG Lemma A11. An equivalent argument holds for h_C^* . ■

In words, Claim A7 says that the object that is clinched at h_B^*/h_C^* is not lurked. Note that the claim also implies that in both Γ_B and Γ_C , $\sigma_B(r_{n'}) = \sigma_C(r_{n'}) = i_{n'}$ for all $n' = 1, \dots, n$, and so $i_{h_B^*} = i$ who clinches y first in Γ_B , and $i_{h_C^*} = j$, who clinches z first in Γ_C .

We can now finish the proof of Lemma A9. Recall that $s = \rho(h_A^*)$ is the role of the first agent to clinch in Γ_A , and s' is the role of the other active non-lurker at h_A^* (for Γ_A , we know that $\sigma_A(s') = j$). There are two cases.

Case (1): $\sigma_A(s) = i$. This is the case where the two non-lurkers at h_A^* are $\{i, j\}$, so that $s = \sigma_A^{-1}(i)$ and $s' = \sigma_A^{-1}(j)$. Note that at h_A^* , both $s, s' \in \bar{\mathcal{L}}_R(h_A^*)$, and $y \in C_{s'}^{\subseteq}(h_A^*)$. Now, consider $\sigma_B/\triangleright_B$. By the discussion following Claim A7, $i_{h_B^*} = i$. This implies that $\sigma_B^{-1}(i) = s'$.⁹²

Now consider $\sigma_C/\triangleright_C$. Again, by the discussion following Claim A7, $i_{h_C^*} = j$; further, $\sigma_C^{-1}(j) = s$.⁹³ If $z \in C_s^{\subseteq}(h_B^*)$, then, by an argument equivalent to footnote 87, it must be that $\sigma_B(s) = j$. Further, $h_B^* \subsetneq h_C^*$ (otherwise, j would clinch at $h_C^* \subsetneq h_B^*$ in Γ_B , since she has the same role in both games). Therefore, $y \in C_{s'}^{\subseteq}(h_C^*)$ (in particular, $y \in C_{s'}(h_B^*)$). Again by the same argument as footnote 87, $\sigma_C(s') = i$. But, $\sigma_C(s') = i$ implies that i clinches at $h_B^* \subsetneq h_C^*$ in Γ_C (since i has the same role as in Γ_B), which is a contradiction. Finally, if $z \notin C_s^{\subseteq}(h_B^*)$, we once again have $h_B^* \subsetneq h_C^*$ and so $y \in C_{s'}^{\subseteq}(h_C^*)$, and we reach the same contradiction as in the previous case.

Case (2): $\sigma_A(s) \neq i$. In this case, in Γ_A , i must be a lurker for some $x_{n'}$, and the first agent to clinch is some $i_{h_A^*} = i_{n_1}$ who clinches some lurked object x_{n_1} . This causes a chain of assignments of lurked objects, that ends with some other agent $i_{n'}$ taking $x_{n'}$, after which i clinches y and all lurked objects $x_{n'+1}, \dots, x_n$ are immediately assigned to their lurkers. Note that $y \in C_{s'}(h_A^*)$ here, by construction of Γ_A . So, we have $\sigma_A^{-1}(s) = i_{h_A^*} \neq i$, and $\sigma_A^{-1}(s') = j$. Since there is a lurked object $x_{n'} \in C_s(h_A^*)$, there is no passing action at h_A^* , by BG Lemma A11. Further, this implies that role s is the terminator role defined in BG Lemma A14.

In game Γ_B , the discussion following Claim A7 again gives $i_{h_B^*} = i$. If $\sigma_B^{-1}(i) = s$, then,

⁹²In particular, we cannot have $\sigma_B^{-1}(i) = s$ again, because this would imply $h_B^* = h_A^*$, contradicting Claim A6.

⁹³If $\sigma_C^{-1}(j) = s'$, then, j must pass at all $h' \subsetneq h_A^*$ at which she is called to play (since $\sigma_A(s') = j$, and j passed at all such h' in Γ_A). Since $i_{h_C^*} = j$, this implies $h_C^* \supsetneq h_A^*$, which contradicts Claim A6.

since i is the first agent to clinch and is the terminator, we have $Top(\succ_i, \mathcal{X}) = y$. But, in game Γ_A , i was a lurker for some $x_{n'} \neq y$ (since y is not a lurked object at h_B^*), which implies $Top(\succ_i, \mathcal{X}) \neq y$, a contradiction. Therefore, $\sigma_B^{-1}(i) = s'$.

In game Γ_C , the discussion following Claim A7 gives $i_{h_C^*} = j$. Just as in Case (1) above, we can show that $\sigma_C^{-1}(j) = s$. Since s is the terminator role of BG Lemma A14, and j clinches z at h_C^* , we conclude that $Top(\succ_j, \mathcal{X}) = z$. If $z \in C_s^C(h_B^*)$, then let $h' \subsetneq h_B^*$ be a history such that $z \in C_s(h')$. By an argument equivalent to footnote 87, we have $\sigma_B(s) = j$. However, this implies that j must clinch z at h' in Γ_B , which contradicts that $i_{h_B^*} = i$. Thus, $z \notin C_s^C(h_B^*)$, which implies that $h_C^* \supsetneq h_B^*$, and so $y \in C_{s'}^C(h_C^*)$ (in particular, $y \in C_{s'}(h_B^*)$). By an argument equivalent to footnote 87, $\sigma_C(s') = i$. However, if $\sigma_C(s') = i$, then i clinches at $h_B^* \subsetneq h_C^*$ in Γ_C (since $\sigma_C(s') = i = \sigma_B(s')$), which contradicts that the first agent to clinch in Γ_C is $i_{h_C^*} = j$ at h_C^* .

We have thus shown that there cannot be three (initial) partial orderings $\triangleright_A, \triangleright_B, \triangleright_C$ of the form given in the statement of Lemma A9 for the first step of the ordering algorithm. For the remaining steps, notice that the subgame after clearing all of the agents in step 1 is simply another millipede game with lurkers (possibly being with one agent from the first stage carrying over her “endowment” to the second). We then simply repeat the same arguments as above, step-by-step, until we reach the end of the game. This completes the proof of Lemma A9.

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.
- ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica*, 64, 51–76.
- ARRIBILLAGA, R. P., J. MASSÓ, AND A. NEME (2017): “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.
- ARTEMOV, G., Y.-K. CHE, AND Y. HE (2017): “Strategic ‘Mistakes’: Implications for Market Design Research,” Tech. rep., mimeo.

- ASHLAGI, I. AND Y. A. GONCZAROWSKI (2018): “Stable matching mechanisms are not obviously strategy-proof,” *Journal of Economic Theory*, 177, 405–425.
- BADE, S. AND Y. GONCZAROWSKI (2017): “Gibbard-Satterthwaite Success Stories and Obvious Strategyproofness,” .
- BO, I. AND R. HAKIMOV (2019): “Pick-an-Object Mechanisms,” Presentation at WZB Designing and Evaluating Matching Markets Conference.
- BOGOMOLNAIA, A. AND H. MOULIN (2001): “A New Solution to the Random Assignment Problem,” *Journal of Economic Theory*, 100, 295–328.
- BÖRGER, T. AND J. LI (2019): “Strategically Simple Mechanisms,” *Econometrica*, *forthcoming*.
- BREITMOSER, Y. AND S. SCHWEIGHOFER-KODRITSCH (2019): “Obviousness Around the Clock,” WZB Discussion Paper SP II 2019–203.
- BUDISH, E. AND E. CANTILLON (2012): “The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard,” *American Economic Review*, 102, 2237–71.
- CARROLL, G. (2014): “A general equivalence theorem for allocation of indivisible objects,” *Journal of Mathematical Economics*, 51, 163–177.
- CHAWLA, S., J. D. HARTLINE, D. L. MALEC, AND B. SIVAN (2010): “Multi-parameter mechanism design and sequential posted pricing,” in *Proceedings of the forty-second ACM symposium on Theory of computing*, ACM, 311–320.
- CHE, Y.-K. AND F. KOJIMA (2010): “Asymptotic Equivalence of Random Priority and Probabilistic Serial Mechanisms,” *Econometrica*, 78, 1625–1672.
- CLARKE, E. H. (1971): “Multipart Pricing of Public Goods,” *Public Choice*, 11, 17–33.
- COPIC, J. AND C. PONSATI (2016): “Optimal Robust Bilateral Trade: Risk Neutrality,” *Journal of Economic Theory*, 163, 276–287.
- DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): “The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility,” *Review of Economic Studies*, 46, 185–216.
- DELACRÉTAZ, D., S. D. KOMINERS, AND A. TEYTELBOYM (2016): “Refugee resettlement,” *University of Oxford Department of Economics Working Paper*.

- 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.
- EHLERS, L. AND T. MORRILL (2017): “(II) Legal Assignments in School Choice,” .
- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2018): “Auctions versus posted prices in online markets,” *Journal of Political Economy*, 126, 178–215.
- FELDMAN, M., N. GRAVIN, AND B. LUCIER (2014): “Combinatorial auctions via posted prices,” in *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, SIAM, 123–135.
- GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On the equivalence of Bayesian and dominant strategy implementation,” *Econometrica*, 81, 197–220.
- 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.
- GREEN, J. AND J.-J. LAFFONT (1977): “Characterization of Satisfactory Mechanisms for Revelation of Preferences for Public Goods,” *Econometrica*, 45, 427–438.
- GROVES, T. (1973): “Incentives in Teams,” *Econometrica*, 41, 617–631.
- GUL, F. AND W. PESENDORFER (2001): “Temptation and Self-Control,” *Econometrica*, 69, 1403–1435.
- (2004): “Self-control and the theory of consumption,” *Econometrica*, 72, 119–158.
- HAGERTY, K. M. AND W. P. ROGERSON (1987): “Robust Trading Mechanisms,” *Journal of Economic Theory*, 42, 94–107.
- HAKIMOV, R. AND O. KESTEN (2014): “The Equitable Top Trading Cycles Mechanism for School Choice,” WZB and Carnegie Mellon University, Working Paper.

- HASSIDIM, A., A. ROMM, AND R. I. SHORRER (2016): “Strategic behavior in a strategy-proof environment,” in *Proceedings of the 2016 ACM Conference on Economics and Computation*, ACM, 763–764.
- HATFIELD, J. W. (2009): “Strategy-Proof, Efficient, and Nonbossy Quota Allocations,” *Social Choice and Welfare*, 33 No. 3, 505–515.
- HOLMSTROM, B. (1979): “Groves’ Scheme on Restricted Domains,” *Econometrica*, 47, 1137–1144.
- JEHIEL, P. (1995): “Limited Horizon Forecast in Repeated Alternate Games,” *Journal of Economic Theory*, 67, 497–519.
- (2001): “Limited Foresight May Force Cooperation,” *The Review of Economic Studies*, 68, 369–391.
- JONES, W. AND A. TEYTELBOYM (2016): “Choices, preferences and priorities in a matching system for refugees,” *Forced Migration Review*.
- KAGEL, J. H., R. M. HARSTAD, AND D. LEVIN (1987): “Information impact and allocation rules in auctions with affiliated private values: A laboratory study,” *Econometrica: Journal of the Econometric Society*, 1275–1304.
- KE, S. (2015): “Boundedly rational backward induction,” Tech. rep., Mimeo, Princeton University.
- KOHLBERG, E. AND J.-F. MERTENS (1986): *Econometrica*, 54, 1003–1037.
- LAIBSON, D. (1997): “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 112, 443–478.
- LI, S. (2017a): “Obvious ex post equilibrium,” *American Economic Review*, 107, 230–34.
- (2017b): “Obviously Strategy-Proof Mechanisms,” *American Economic Review*, 107, 3257–87.
- LIU, Q. AND M. PYCIA (2011): “Ordinal Efficiency, Fairness, and Incentives in Large Markets,” .
- LOERTSCHER, S. AND L. M. MARX (2015): “Prior-Free Bayesian Optimal Double-Clock Auctions,” Working Paper.

- MACKENZIE, A. (2017): “A Revelation Principle for Obviously Strategy-proof Implementation,” *working paper*.
- MANELLI, A. M. AND D. R. VINCENT (2010): “Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model,” *Econometrica*, 78, 1905–1938.
- MORRILL, T. (2014): “Making Just School Assignments,” *Games and Economic Behavior*, (forthcoming).
- MYERSON, R. B. (1981): “Optimal auction design,” *Mathematics of operations research*, 6, 58–73.
- 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 (2011): “Lotteries in student assignment: An equivalence result,” *Theoretical Economics*, 6, 1–17.
- PYCIA, M. (2011): “Ordinal Efficiency, Fairness, and Incentives in Large Multi-Unit-Demand Assignments,” Working Paper.
- (2012): “Stability and Preference Alignment in Matching and Coalition Formation,” *Econometrica*, 80, 323–362.
- (2017): “Invariance and Matching Market Outcomes,” Working Paper.
- PYCIA, M. AND M. U. ÜNVER (2016): “Arrovian Efficiency in Allocation of Discrete Resources,” UCLA and Boston College, Working Paper.
- (2017): “Incentive Compatible Allocation and Exchange of Discrete Resources,” *Theoretical Economics*, 12, 287–329.
- REES-JONES, A. (2017): “Mistaken play in the deferred acceptance algorithm: Implications for positive assortative matching,” *American Economic Review*, 107, 225–29.
- (2018): “Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match,” *Games and Economic Behavior*, 108, 317–330.
- RILEY, J. G. AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71, 381–392.

- ROSENTHAL, R. W. (1981): “Games of perfect information, predatory pricing and the chain-store paradox,” *Journal of Economic theory*, 25, 92–100.
- 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.
- SAVAGE, L. J. (1954): *Foundations of Statistics*, New York: John Wiley and Sons.
- SHAPLEY, L. AND H. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 23–37.
- SHORRER, R. I. AND S. SÓVÁGÓ (2018): “Obvious mistakes in a strategically simple college admissions environment: Causes and consequences,” *Available at SSRN 2993538*.
- SÖNMEZ, T. AND M. U. ÜNVER (2010): “Course Bidding at Business Schools,” *International Economic Review*, 51, 99–123.
- TROYAN, P. (2019): “Obviously Strategy-Proof Implementation of Top Trading Cycles,” *International Economic Review*, 60.
- TROYAN, P., D. DELACRETAZ, AND A. KLOOSTERMAN (2018): “Essentially Stable Matchings,” *working paper*.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- ZHANG, L. AND D. LEVIN (2017a): “Bounded Rationality and Robust Mechanism Design: An Axiomatic Approach,” *American Economic Review Papers and Proceedings*, 107, 235–39.
- (2017b): “Partition Obvious Preference and Mechanism Design: Theory and Experiment,” .