

Pycia, Marek; Troyan, Peter

Working Paper

A theory of simplicity in games and mechanism design

Working Paper, No. 393

Provided in Cooperation with:

Department of Economics, University of Zurich

Suggested Citation: Pycia, Marek; Troyan, Peter (2021) : A theory of simplicity in games and mechanism design, Working Paper, No. 393, University of Zurich, Department of Economics, Zurich,
<https://doi.org/10.5167/uzh-204780>

This Version is available at:

<https://hdl.handle.net/10419/235606>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series

ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 393

A Theory of Simplicity in Games and Mechanism Design

Marek Pycia and Peter Troyan

June 2021

A Theory of Simplicity in Games and Mechanism Design

Marek Pycia and Peter Troyan*

June 2021

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 graduated approach to simplicity to provide characterizations of simple mechanisms. While more demanding simplicity standards may reduce the flexibility of the designer in some cases, this is not always true, and many well-known mechanisms are simple, including ascending auctions, posted prices, and serial dictatorship-style mechanisms. In particular, 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, Tilman Börgers, Ernst Fehr, Ben Golub, Yannai Gonczarowski, Ed Green, Rustamdjan Hakimov, Shaowei Ke, Fuhito Kojima, Jiangtao Li, 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, NEG-T’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, Wash U, Maryland, Warsaw Economic Seminars, ISI Delhi, Notre Dame, UCSD, Columbia, Rochester, Brown, Glasgow, Singapore Management University, Matching in Practice, Essex, 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 standard approach in mechanism design is to provide agents with incentives to play in a predictable optimal way. For instance, the designer may use a Bayesian or dominant-strategy incentive-compatible direct mechanism where it is optimal for agents to simply report all of their private information truthfully. This approach succeeds if the participants understand that being truthful is in their interest, but 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, may not be sufficiently simple for participants to play optimally in practice. Simpler mechanisms are also appealing because they lower participation costs, attract participants, and equalize opportunities across participants with different levels of access to information and strategic sophistication. Additionally, designing simpler mechanisms requires less information about participants’ beliefs.²

What mechanisms, then, are actually “simple to play”? We address this question by introducing a general class of simplicity standards that vary the foresight abilities required of agents in extensive-form imperfect-information games. We then use these standards to assess the restrictions simplicity imposes on the mechanism designer, as well as to characterize simple mechanisms for a broad range of social choice environments both with and without transfers.³ Similarly to how the revelation principle allows a designer to limit the search for a Bayesian mechanism to the space of incentive-compatible direct mechanisms, our results construct classes of mechanisms that allows one to do the same when searching for simple mechanisms.⁴ As applications, we provide microfoundations for popular simple mechanisms such as posted prices, ascending auctions, and Random Priority.

¹See, e.g., Kagel, 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).

²See Vickrey (1961) for participation costs, Spenner and Freeman (2012) for attracting participants, Pathak and Sönmez (2008) for leveling the playing field, and Wilson (1987) and Bergemann and Morris (2005) for a designer’s informational requirements.

³Examples include auctions (Vickrey, 1961; Riley and Samuelson, 1981; Myerson, 1981), 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).

⁴Direct mechanisms are in general not simple, and so the revelation principle does not extend to simple extensive form games, cf. Li (2017b).

The main innovation in our approach is a departure from the standard assumption that agents plan a complete strategy for every possible future contingency; rather, we consider agents that, at each information set, make plans for only those information sets that they perceive as simple from the current perspective. We refer to these plans as *partial strategic plans*.⁵ A (partial) strategic plan is *simply dominant* if the called for action is weakly better than any alternative, irrespective of what happens at information sets that are not simple. As the game progresses, the agent’s perception of which information sets are simple may change and we allow agents to update their strategic plans along the path of the game, which is what differentiates strategic plans from the standard game-theoretic concept of a strategy.⁶

Variations in the sets of information sets perceived as simple gives rise to a family of simple dominance standards that vary in strength. The stronger the simplicity standard—i.e., the fewer information sets in the future that are perceived as simple from today’s perspective—the more robust the corresponding mechanisms will be to agents who can plan for only limited future horizons (e.g., because of non-exponential discounting) or whose decision capabilities are otherwise constrained.⁷ We focus on special cases of simple dominance in which agents are able to plan some exogenously given number $k \in \{0, 1, \dots, \infty\}$ of future moves; that is, they perceive as simple their current information set and only the first k information sets at which they may be called to play in the continuation game. We show that the longer the foresight horizon of the agents, the more social choice rules a designer can implement in a simply-dominant way; furthermore, without loss of flexibility, the designer can restrict attention to perfect-information extensive-form games.

We analyze three special cases of simple dominance in detail.

- $k = \infty$: 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, an agent can plan the actions they will take at any future information set at which they may be called to play. This is equivalent to Li’s (2017b) notion of obvious dominance; for this

⁵Savage (1954) wrestles with whether decision-makers should be modeled as “look before you leap” (create a complete contingent 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.

⁶We are agnostic as to whether the agents are sophisticated and understand that their plans might be updated, or whether the agents are naive about this possibility. Simple dominance only requires that the initial action of the strategic plan is unambiguously better than other actions the agent could have chosen at the information set at which the plan is made; the subsequent actions of the strategic plan merely ensure the optimality of the initial action.

⁷We show that a strategic plan is simply dominant if and only if in every game an agent may confuse with the actual game being played, the strategic plan is weakly dominant in the standard sense (Theorem 3). Li (2017b) provides a related behavioral microfoundation for his obvious dominance, on which we build.

reason, we refer to the resulting simply dominant strategic plans as *obviously dominant*, and the corresponding mechanisms as *obviously strategy-proof (OSP)*.

- $k = 1$: agents 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. We refer to the resulting simply dominant strategic plans as *one-step dominant*, and the corresponding mechanisms as *one-step simple (OSS)*.
- $k = 0$: agents perceive as simple *only* their own current information set—in other words, agents cannot plan for any moves in the future. We refer to the resulting simply dominant strategic plans as *strongly obviously dominant*, and the corresponding mechanisms as *strongly obviously strategy-proof (SOSP)*.

The above concepts are nested: strongly obviously dominant strategic plans are one-step dominant, which in turn are obviously dominant. While obvious dominance is the most permissive standard, it relies on the assumption that agents can create a complete plan for all possible contingencies going forward, and further are able to perform backwards induction over at least their own future actions (though not over the actions of their opponents). As an example, consider the game of 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. On the other hand, winning strategies in chess are not one-step dominant, nor strongly obviously dominant, as they require looking many steps into the future. Games that admit one-step and/or strongly obviously dominant strategies do not require agents to have such lengthy foresight.

For the above three simplicity standards we ask: which mechanisms are simple? For obvious dominance, we focus on social choice environments without transfers, hence complementing Li (2017b), who focuses on the case with transfers. We show that OSP games can be represented as *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 and leaves the game, and the process continues with the next agent. 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 (see Figure 2 in Section 4.2 for an example of such a game).

We next study one-step dominance in environments both with and without transfers. We first show that in the binary allocation environments with transfers studied by Li (2017b)—which encompass canonical special cases such as single-unit auctions and binary public good choice—any one-step simple mechanism is equivalent to a personal clock auction. This strengthens Li’s result that personal clock auctions are OSP by showing that there is no loss in imposing OSS relative to OSP: any social choice rule that is implementable in obviously dominant strategies is also implementable in one-step dominant strategic plans. In no-transfer environments, one-step simplicity eliminates the complex, yet still formally, OSP millipede games discussed above (and also eliminates games such as chess). Indeed, we can characterize OSS millipede games as those that satisfy the following monotonicity property: 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. Monotonic games seem particularly simple, both for a designer to implement, since the agent only needs to recognize that she can do no worse at her very next move if she remains in the game.⁸

For strong obvious dominance, we show that SOSDP games do not require agents to look far into the future and perform lengthy backwards induction: in all such games, each agent has essentially at most one payoff-relevant move. Thus, strongly obviously dominant strategic plans are robust to agents who may be concerned about trembles, or have time-inconsistent preferences. Building on this insight, we show that all SOSDP games can be implemented as sequential choice games in which each agent moves at most once, and, at this move, is offered a choice from a menu of options. 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. The offered menu may include prices, in which case we call the mechanism a *(sequential) posted price mechanism*. In this way, strong obvious dominance gives us a microfoundation for posted prices, a ubiquitous sales mechanism.⁹

As an application of our analysis, we provide an axiomatic characterization of the well-

⁸Note also that our monotonicity property is a generalization of a similar feature of ascending auctions (and also Li’s personal clock auctions). In an ascending auction, if an agent passes (continues in the auction), at any next move, she will be offered the opportunity to drop out (clinch the zero payoff), except if she wins.

⁹For earlier microfoundations of posted prices, see Hagerty and Rogerson (1987) and Copic and Ponsati (2016).

known Random Priority (RP; also known as Random Serial Dictatorship) mechanism using simplicity, efficiency, and fairness axioms. 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, and 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 (as well as one-step simple and strongly obviously strategy-proof). 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.¹⁰ 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 construction of the simplicity criteria is inspired by Li (2017b), who formalized obvious strategy-proofness and established its desirability as an incentive property; we go beyond his work by allowing for the gradation of simplicity criteria—which allows us to assess the trade-off between simplicity and implementation flexibility—and by providing simplicity-based microfoundations for popular mechanisms such as posted prices and Random Priority. Following up on Li’s work, but preceding ours, Ashlagi and Gonczarowski (2018) show that stable mechanisms such as Deferred Acceptance (DA) are not obviously strategy-proof, except in very restrictive environments where DA 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.¹¹ Following our work, Arribillaga et al. (2017) characterize the voting rules that are obviously

¹⁰The efficiency and fairness of Random Priority were recognized already by Abdulkadiroğlu and Sönmez (1998), while its obvious strategy-proofness was established by Li (2017b). In single-unit demand allocation with at most three agents and three objects, Bogomolnaia and Moulin (2001) proved that Random Priority is the unique mechanism that is strategy-proof, efficient, and symmetric. In markets in which each object is represented by many copies, Liu and Pycia (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.

¹¹Li 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. Following our and Troyan’s work, Mandal and Roy (2020) characterize the class of OSP, efficient, and non-bossy social choice rules in this setting.

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.¹²

Our work contributes to the understanding of limited foresight and limits on backward induction. Other work in this area—with different approaches from ours—includes Jehiel (1995; 2001) on limited foresight equilibrium in which players’ forecasts are correct, Gabaix et al. (2006) on directed cognition, Ke’s (2019) axiomatization of bounded-horizon backward induction, as well as the rich literature on time-inconsistent preferences (e.g., Laibson (1997) and Gul and Pesendorfer (2001; 2004)). A major difficulty for models of imperfect foresight is the question of how an agent takes into account the future they are unable to foresee; we resolve this difficulty by designing games in which all resolutions of the unforeseen lead the agent to the same current decision.¹³

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 (2020; 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.¹⁴ Liu and Pycia (2011), Pycia (2011), Morrill (2014), Hakimov and Kesten (2014), Ehlers and Morrill (2017), and Troyan et al. (2020) characterize mechanisms that satisfy incentive, efficiency, and fairness objectives.

¹²Also of note is Loertscher and Marx (2020) 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. A different strategic perspective on simplicity in mechanism design was explored by Börgers and Li (2019).

¹³The issue of accounting for the unforeseen is also crucial for the analyses of incomplete contracts (e.g., Maskin and Tirole, 1999) and unawareness (e.g., Karni and Viero, 2013). Agents who rely on incomplete models have been also studied in the context of persuasion (e.g., Schwartzstein and Sunderam, 2021).

¹⁴Pycia and Ünver (2020) characterized individually strategy-proof and Arrovian efficient mechanisms. For an analysis of these issues under additional feasibility constraints, see also Dur and Ünver (2015).

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.¹⁵ An outcome might involve a monetary transfer; we allow both environments with and without transfers. Each agent has a preference ranking over outcomes, where, for any two $x, y \in \mathcal{X}$, we write $x \succeq_i y$ to denote that x is weakly preferred to y . We allow for indifferences, and write $x \sim_i y$ if $x \succeq_i y$ and $y \succeq_i x$. For any \succeq_i , we let $>_i$ denote the corresponding strict preference relation, i.e., $x >_i y$ if $x \succeq_i y$ but not $y \succeq_i x$. We use \mathcal{P}_i to denote the domain of agent i 's preferences, and will often refer to \succeq_i as agent i 's **type**.

We allow incomplete information through the standard imperfect-information construction of a meta-game in which Nature moves first and determines agents' types, and only then the designed game/mechanism is played. Due to the nature of the dominance properties we study, we do not need to make any assumptions on agents' beliefs about others' types nor on how agents' evaluate lotteries.¹⁶

2.2 Extensive Form Games

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**, \mathcal{H} . We write $h' \subseteq h$ to denote that $h' \in \mathcal{H}$ is a subhistory of $h \in \mathcal{H}$, and $h' \subset h$ when $h' \subseteq h$ but $h \neq h'$. Terminal histories will be denoted with bars, i.e., \bar{h} . Each $\bar{h} \in \mathcal{H}$ is associated with an outcome in \mathcal{X} . 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.

¹⁵Assuming \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.

¹⁶It is natural to assume that an agent weakly prefers lottery μ over ν whenever for all outcomes $x \in \text{supp}(\mu)$ and $y \in \text{supp}(\nu)$ this agent weakly prefers x over y . This mild assumption is satisfied for expected utility agents, as well as for agents who prefer μ to ν as soon as μ first-order stochastically dominates ν . While our results do not rely on this assumption, it ensures that dominant actions always lead to weakly preferred lotteries over outcomes.

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 \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 . We say that an information set I_1 **precedes** information set I_2 if there are $h_1 \in I_1$ and $h_2 \in I_2$ such that $h_1 \subseteq h_2$; we then write $I_1 \leq I_2$ (and $I_1 < I_2$ if $I_1 \neq I_2$) and we also say that I_2 **follows** I_1 and that I_2 is a **continuation** of I_1 . We say that an outcome x is **possible** at information set I if there is $h \in I$ and a terminal history $\bar{h} \supseteq h$ such that x obtains at \bar{h} .

3 Simple Dominance

We propose a class of simplicity standards that relaxes the standard assumption of economic analysis that players can analyze and plan their actions arbitrarily far into the future of the game. Such foresight assumptions are embedded in standard game theoretic concepts of backward induction, dynamic programming, perfect Bayesian equilibrium, iterated dominance, weak dominance, and Li’s obvious dominance. In relaxing the foresight assumption, we build on the pioneering approach of Li (2017b) whose obvious dominance allows for players who cannot fully analyze the actions of others, but maintains the assumption that players understand the set of possible outcomes following their own actions and the structure of precedence among their own information sets. While obvious dominance 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. It assumes that they know all of *their own* actions they will take in the future, and understand precisely the sets of outcomes that could possibly obtain, conditional on any sequence of their own actions that they plan to take. This is the reason that if White has a winning strategy in chess then—knowing at the start of the game what she needs to do at any possible future configuration of the board in order to ensure a victory—White has a strategy that is not only winning, but also obviously dominant. We relax Li’s foresight assumptions, only maintaining that players know possible outcomes of actions at information sets they perceive as simple and the precedence relations among these information sets.

The key innovation in our framework is that the information sets an agent perceives as simple may update as the game is played. In other words, we 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” as a *complete* contingent plan of action—to vary as the game progresses.

Formally, for each player i and information set $I^* \in \mathcal{I}_i$ at which i moves, there is a set of information sets $\mathcal{I}_{i,I^*} \subseteq \{I \in \mathcal{I}_i | I \supseteq I^*\}$ that are perceived as **simple information sets** from the perspective of I^* . We assume that $I^* \in \mathcal{I}_{i,I^*}$, but otherwise, the only restriction is that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}_i$.¹⁷ A **(partial) strategic plan** $S_{i,I^*}(>_i)$ for agent i of type $>_i$ at information set I^* maps each simple information set $I \in \mathcal{I}_{i,I^*}$ to an action at this information set.¹⁸ Note that a strategic plan does not specify the play at all continuation information sets at which i may be called to move, but rather only at the information sets that are simple from the perspective of I^* . Sets of strategic plans $(S_{i,I^*}(>_i))_{I^* \in \mathcal{I}_i}$ and $(S_{i,I^*}(>_i))_{I^* \in \mathcal{I}_i, >_i \in \mathcal{P}_i}$ of agent i are called **strategic collections**.

An **extensive-form mechanism** $(\Gamma, S_{\mathcal{N},\mathcal{I}})$, or simply a **mechanism**, is an extensive-form game Γ together with a profile of strategic collections, $S_{\mathcal{N},\mathcal{I}} = ((S_{i,I^*}(>_i))_{I^* \in \mathcal{I}_i, >_i \in \mathcal{P}_i})_{i \in \mathcal{N}}$. For any strategic collection $(S_{i,I^*}(>_i))_{I^* \in \mathcal{I}_i}$, we define the **induced strategy** $\hat{S}_i(>_i) : \mathcal{I}_i \rightarrow \cup_{I \in \mathcal{I}_i} A(I)$ as the mapping from information sets to actions defined by $\hat{S}_i(>_i)(I) = S_{i,I}(>_i)(I)$ for each $I \in \mathcal{I}_i$; that is, $\hat{S}_i(>_i)$ is a standard game-theoretic strategy (complete contingent plan of action) defined by agent i selecting the action that is called for by the strategic plan $S_{i,I}$ at information set I itself. For any $S_{\mathcal{N},\mathcal{I}}$ and type realization $>_{\mathcal{N}}$, we can find the terminal history/outcome that is reached when the game is played according to the profile of strategic collections $S_{\mathcal{N},\mathcal{I}}(>_{\mathcal{N}})$ by following the profile of induced strategies $\hat{S}_{\mathcal{N}}(>_{\mathcal{N}})$. For each player i and type $>_i$, the induced strategy $\hat{S}_i(>_i)$ allows us to define the set of **on-path information sets** for a strategic collection as the information sets $I \in \mathcal{I}_i$ such that there exists other players’ and Nature’s strategies such that I is on the path of play of $\hat{S}_i(>_i)$.

Induced strategies also allow us to define equivalence of mechanisms: two mechanisms $(\Gamma, S_{\mathcal{N},\mathcal{I}})$ and $(\Gamma', S'_{\mathcal{N},\mathcal{I}})$ are **equivalent** if, for every profile of types $>_{\mathcal{N}}$, the distribution over outcomes from the induced strategies $\hat{S}_{\mathcal{N}}(>_{\mathcal{N}})$ in Γ is the same as from the induced strategies $\hat{S}'_{\mathcal{N}}(>_{\mathcal{N}})$ in $\hat{\Gamma}$. This equivalence definition is purely outcome-based, and allows that $(\Gamma, S_{\mathcal{N},\mathcal{I}})$ and $(\Gamma', S'_{\mathcal{N},\mathcal{I}})$ have different classes of simple information sets. Given a mechanism, we can construct the corresponding **social choice rule**—that is, mapping from preference profiles to outcomes—that is implemented. All mechanisms in the same equivalence class implement

¹⁷The assumption that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}_i$ is made for simplicity; in its absence we need to endow players with beliefs of what other players will do. A natural requirement on the collection of simple node sets is that if an agent classifies an information set $I > I_1$ as simple from the perspective of information set I_1 then the agent continues to classify I as simple from the perspective of all information sets $I_2 > I_1$ such that $I \geq I_2$; while we do not impose this requirement, it is satisfied in all of the examples of simple dominance that we study.

¹⁸We focus on pure strategies; the extension to mixed strategies is straightforward.

the same social choice rule.

Strategic plan $S_{i,I^*}(>_i)$ is **simply dominant at information set I^*** for type $>_i$ of player i if the worst possible outcome for i in the continuation game assuming i follows $S_{i,I^*}(>_i)(I)$ at all $I \in \mathcal{I}_{i,I^*}$ is weakly preferred by i to the best possible outcome for i in the continuation game if i plays some other action $a' \neq S_{i,I^*}(>_i)(I^*)$ at I^* . We say that a strategic collection $(S_{i,I^*}(>_i))_{I^* \in \mathcal{I}_i, >_i \in \mathcal{P}_i}$ is **simply dominant** if, for each type $>_i \in \mathcal{P}_i$, the strategic plan $S_{i,I^*}(>_i)$ is simply dominant at I^* for each on-path information set I^* .¹⁹ We say that a game is simple dominant if it admits simply dominant strategies.

Note that the collections of simple information sets, $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$, is a parameter of the model. In the sequel, we focus on collections of simple information sets that vary in the foresight of the agents, though this is not necessary, and there are other ways to conceptualize what information sets are viewed as simple from a given perspective.²⁰ Given a fixed $k \in \{0, 1, 2, \dots, \infty\}$, we say that agent i can plan k moves ahead but not more if $\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I \text{ and } I^* < I_1 \dots < I_k < I \Rightarrow \exists \ell \in \{1, \dots, k\} \text{ s.t. } I_\ell \notin \mathcal{I}_i\}$. We refer to the resulting simply dominant strategic collections as **k -step dominant** and we say that a strategy is k -simple if it is the induced strategy for some k -step dominant strategic collection. Varying k allows us to embed in our model the following special cases:

- $k = \infty$ that is $\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I\}$; i can plan all of her future moves. We refer to the resulting simply dominant strategic collections as **obviously dominant**, because the induced strategy is obviously dominant in the sense of Li (2017b), and any obviously dominant strategy S_i in the sense of Li (2017b) determines an obviously dominant strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ by defining $S_{i,I^*}(I) = S_i(I)$ for any $I^* \leq I$. If a mechanism admits obviously dominant strategic collections, then we say it is **obviously strategy-proof (OSP)**.
- $k = 1$ that is $\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I \text{ and } I^* < I' < I \Rightarrow I' \notin \mathcal{I}_i\}$; i can plan one move ahead but not more. We refer to the resulting simply dominant strategic collections as **one-step dominant**. The information sets in $\mathcal{I}_{i,I^*} - \{I^*\}$ are called i 's **next information**

¹⁹When assessing $S_{i,I^*}(>_i)(I)$, we take the worst case over all game paths consistent with i following $S_{i,I^*}(>_i)(I)$ at all $I \in \mathcal{I}_{i,I^*}$, and compare to the best case over *all* game paths following any alternative action $a' \neq S_{i,I^*}(>_i)(I^*)$. While formulated slightly differently than Li (2017b), who invokes the notion of an earliest point of departure between two strategies, our definition is formally equivalent to his when \mathcal{I}_{i,I^*} is the set of all continuation information sets at which i moves. Both we and Li (2017b) require simple dominance (respectively, obvious dominance) only on-path; this choice is in line with e.g. Pearce's (1984) extensive form rationalizability and Shimoji and Watson's (1998) conditional dominance. An alternative approach is to require simple dominance at all nodes (information sets) in the game, including off-path ones.

²⁰For instance, the collection of simple information sets could be those at which a measure of computational complexity of a decision problem is below some threshold; cf. e.g. Arora and Barak (2009) for a survey of computational complexity criteria.

sets (from the perspective of I^*). If a mechanism admits one-step dominant strategic collections, then we say it is **one-step simple (OSS)**.

- $k = 0$ that is $\mathcal{I}_{i,I^*} = \{I^*\}$; i cannot plan any future moves. We refer to the resulting simply dominant strategic collections as **strongly obviously dominant**. In this case, we can also talk about strongly obviously dominant *strategies* because, as for obvious dominance, there is a one-to-one correspondence between strategic collections $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ and the induced strategies $\hat{S}_i(I^*) = S_{i,I^*}(I^*)$. If a mechanism admits strongly obviously dominant strategic collections, then we say it is **strongly obviously strategy-proof (SOSP)**.

Remark 1. Obviously dominant strategic collections and strongly obviously dominant strategic collections are consistent in the following sense: $S_{i,I^*}(I) = S_{i,I}(I)$ for all $I \in \mathcal{I}_{i,I^*}$ and all $I^* \in \mathcal{I}_i$. One-step dominant strategic collections, on the other hand, do not satisfy this property; ascending auctions, discussed in Section 4.2, are an example of such a case. The failure of this property does not mean that agents who plan only one step ahead are inconsistent or irrational. Indeed, such agents might understand that they may adjust their plans later, and think of the partial strategic plan S_{i,I^*} as an argument establishing that playing $S_{i,I^*}(I^*)$ is better than any other action they could take at I^* . The tentativeness of such partial plans is an important possibility in the under-explored game-theoretic paradigm of making choices as they arise, a paradigm that Savage (1954) describes as “you can cross that bridge when you come to it” (cf. Introduction).

A direct verification shows that the smaller the set of simple nodes, the stronger is the resulting simplicity requirement. To formulate this result, for any simple information sets \mathcal{I}_{i,I^*} and \mathcal{I}'_{i,I^*} such that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}'_{i,I^*}$, we say that a strategic collection $(S'_{i,I^*}(>_i))_{>_i \in \mathcal{P}_i}$ on \mathcal{I}'_{i,I^*} is an \mathcal{I}'_{i,I^*} -extension of a strategic collection $(S_{i,I^*}(>_i))_{>_i \in \mathcal{P}_i}$ on \mathcal{I}_{i,I^*} if $S'_{i,I^*}(>_i)(I) = S_{i,I^*}(>_i)(I)$ for all $I \in \mathcal{I}_{i,I^*}$.

Theorem 1. (Nesting of Simplicity Concepts). *If simple information sets \mathcal{I}_{i,I^*} and \mathcal{I}'_{i,I^*} are such that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}'_{i,I^*}$ and strategic collection S_{i,I^*} is simply dominant at I^* for \mathcal{I}_{i,I^*} , then any \mathcal{I}'_{i,I^*} -extension of S_{i,I^*} is simply dominant at I^* for \mathcal{I}'_{i,I^*} .*

As a corollary, we conclude that the lower the parameter k , the more restrictive k -step simplicity becomes. Further, our class of simple dominance concepts has a natural lattice structure, with obvious dominance as its least demanding concept and strong obvious dominance as the most demanding one.

Corollary 1. *(i) Take $k, k' \in \{0, 1, 2, \dots, \infty\}$ and assume $k < k'$. Then, any strategic collection that is k -step dominant is also k' -step dominant.*

(ii) If a strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ is simply dominant for some collection of simple information sets, then the induced strategy $\hat{S}_i(I^*) = S_{i,I^*}(I^*)$ is obviously dominant.

(iii) If the induced strategy $\hat{S}_i(I^*) = S_{i,I^*}(I^*)$ is strongly obviously dominant, then the strategic collection is simply dominant for any $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$.

From an implementation perspective, an immediate consequence of Corollary 1 is that the set of k -step simple implementable social choice rules weakly expands as k is increased. The following result shows that in general, this inclusion is strict: that is, stronger simplicity constraints (lower k) reduce the flexibility of the designer.²¹

Theorem 2. *Let $k, k' \in \{0, 1, 2, \dots, \infty\}$ and assume $k' > k$. There exist social choice rules that are implementable in k' -step simple strategic collections, but are not implementable in k -step simple strategic collections.*

The presence of the simplicity-implementability trade-off depends on the preference environment: for instance, Theorem 6 shows that in some environments there is no loss in imposing one-step simplicity ($k = 1$) relative to obvious strategy-proofness ($k = \infty$): in these environments, any social choice rule that is OSP-implementable is also OSS-implementable.

To get a sense of why the inclusion can be strict consider an environment with transfers in which there are at least two agents and each agents' values come from the same support with at least three distinct values. Suppose we want to allocate an object to the highest-value agent. This social choice rule can be implemented via an ascending auction and ascending auctions are OSS (we establish the one-step simplicity of ascending auctions in Theorem 6). At the same time, this social choice rule, and the price discovery it entails, cannot be implemented via SOSp mechanisms, which resemble posted prices (the posted price characterization of SOSp is given by our Theorem 8). For k, k' strictly larger than 0, our proof in the appendix constructs social rules that are k' -step simple implementable but not k -step simple implementable in no-transfer single-unit demand allocation environments.

3.1 Behavioral Microfoundations

We may think of simple strategic plans as providing guidance to a player that is unaffected even when they may be confused about the game they are playing, in the sense that they may mistake the game for a different game that has different players, actions, and precedence relations at non-simple information sets. An alternative interpretation is that the player is only given a partial description of the game: each time they are called to move, they are

²¹In particular, the theorem shows that for any $k < \infty$, there are social choice rules that are OSP-implementable but not k -step implementable.

told what happens at their own simple information sets, but not at any other non-simple information set. If players have simply dominant strategic plans, the prediction of play will be unaffected by concerns that they might be so-confused.

To formalize this idea, say that game Γ' is **indistinguishable from Γ from the perspective of agent i at information set I^* of game Γ** if there is an injection λ from the set of agent i 's simple information sets \mathcal{I}_{i,I^*} in Γ into the set of agent i 's information sets \mathcal{I}'_i in Γ' such that:

1. If $I_1, I_2 \in \mathcal{I}_{i,I^*}$ and I_1 precedes I_2 in Γ , then $\lambda(I_1)$ precedes $\lambda(I_2)$ in Γ' .
2. For each $I \in \mathcal{I}_{i,I^*}$, there is a bijection η_I that maps actions at agent's i information set I in Γ onto actions at agent's i information set $\lambda(I)$ in Γ' .
3. An outcome is possible following action a at $I \in \mathcal{I}_{i,I^*}$ in Γ if and only if this outcome is possible following $\eta_I(a)$ at $\lambda(I)$ in Γ' .

We say that $\lambda(I)$ is the game Γ' **counterpart** of information set I and $\eta_I(a)$ is the game Γ' **counterpart** of action a at information set I in game Γ . The concept of indistinguishability captures the idea that agent i understands the precedence relation among simple information sets, as well as the available actions and possible outcomes at these information sets.

Simple dominance is equivalent the standard weak dominance on all games that are indistinguishable from the game played. We say that a strategy S_i of player i **weakly dominates** strategy S'_i in the continuation game beginning at I^* 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 3. (Behavioral Microfoundation). *For each game Γ , agent i , type $>_i$, and collection of simple information sets $(I_{i,I^*})_{I^* \in \mathcal{I}_i}$, the strategic plan S_{i,I^*} is simply dominant from the perspective of $I^* \in \mathcal{I}_i$ in Γ if and only if, in every game Γ' that is indistinguishable from Γ from the perspective of i at information set I^* , in the continuation game of Γ' starting at the counterpart of I^* , any strategy that at the counterpart of each $I \in \mathcal{I}_{i,I^*}$ selects the counterpart of $S_{i,I^*}(I)$ weakly dominates any strategy that does not select the counterpart of $S_{i,I^*}(I^*)$ at the counterpart of I^* .*

This theorem tells us the strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ is simply dominant in Γ if and only if for every $I^* \in \mathcal{I}_i$ in every game Γ' that is indistinguishable from the perspective of information set I^* every strategy S'_i that agrees with the counterpart of S_i at all $I \in \mathcal{I}_{i,I^*}$ is weakly dominant in the continuation game starting at the counterpart of I^* . When the

strategic collection is consistent, we can express this result equivalently in terms of simplicity of the induced global strategies $S_i(I) = S_{i,I}(I)$. When expressed in this way, this result corresponds to Li’s (2017b) microfoundation for obvious strategy-proofness.²²

3.2 Design Sufficiency of Perfect Information Games

Under perfect information, each information set I contains a single history (or node) h and, to keep the notation at the minimum, we will identify history h and information set $\{h\}$. The key parameter of the simplicity definition then becomes the collection $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$ of **simple histories**, and we denote the corresponding strategic collections by $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$.

Perfect information games plays a special role in designing simply dominant mechanisms because for any imperfect-information simply dominant mechanism, we can find an equivalent perfect-information one.²³ To make this point precise, for any imperfect-information game Γ , define the corresponding perfect information game Γ' with the same set of histories as Γ . Given a collection of simple information sets $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$ in Γ , we define the induced collection of simple histories $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$ in Γ' such that \mathcal{H}_{i,h^*} consists of all histories in \mathcal{I}_{i,I^*} . For a strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$, we define the induced strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ such that $S_{i,h^*}(h) = S_{i,I^*}(I)$, where I is a continuation information set of I^* , $h^* \in I^*$ and $h \in I$.

Theorem 4. (*Perfect-Information Reduction*). *If $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ is simply dominant in an imperfect-information game Γ with simple information sets $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$, then in the corresponding perfect information game Γ' with the induced simple histories $(\mathcal{H}_{i,h^*})_{h^* \in \mathcal{H}_i}$, the induced strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is simply dominant.*

To prove the theorem, consider an agent i with type \succ_i . Notice that if some history h is on-path for the strategic collection $(\hat{S}_{i,h^*}(\succ_i))_{h^* \in \mathcal{H}_i}$ in Γ' , then the corresponding information set $I \ni h$ is on-path for the strategic collection $(\hat{S}_{i,I^*}(\succ_i))_{I^* \in \mathcal{I}_i}$ in Γ . Furthermore, the worst outcome following $S_{i,h^*}(h) = S_{i,I^*}(I)$ in Γ' is weakly better than the worst outcome over the entire information set I when following this strategy. Similarly, the best outcome following an alternative action $a \neq S_{i,h^*}(h)$ at h is worse than the best outcome following an alternative action $a \neq S_{i,I^*}(I)$ over the entire information set I . Thus, if the strategic plan $S_{i,I^*}(I)$ is simply dominant in Γ , then the induced strategic plan $S_{i,h^*}(\succ_i)$ is simply dominant in Γ' .

²²While the two results capture the same phenomenon, there is a slight difference between them even when restricted to OSP, as Li’s (2017b) microfoundation assumes that λ is a bijection on all agent i ’s information sets. We could embed this assumption in our analysis by assuming that λ is a bijection from agent i ’s simple information sets in Γ to simple information sets in Γ' .

²³An analogous property of obvious strategy-proofness was asserted in Ashlagi and Gonczarowski (2018). Following our work, Mackenzie (2017) extended this property of obvious strategy-proofness to extensive-form games without perfect recall.

In light of Theorem 4 we focus on perfect information games in the study of design of the next two sections.

4 Characterizing Simple Mechanisms

We now consider three special cases of the above simplicity standards—obvious dominance, one-step dominance, and strong obvious dominance—and characterize simple mechanisms and social rules in environments both with and without transfers. To make our analysis relevant for market design applications and to avoid general impossibility results such as the Gibbard-Satterthwaite Theorem, we must impose certain assumptions on the domains of agent preferences. We formalize this as follows: We take as a primitive a **structural dominance relation** over outcomes, denoted \succeq , where \succeq is a reflexive and transitive binary relation on \mathcal{X} . The notation $x \succeq y$ is read as “ x weakly dominates y ” or “ x trumps y ”. If $x \succeq y$ but not $y \succeq x$, then we write $x \succ y$, and say that x strictly dominates (or strictly trumps) y . For instance, in environments with transfers, outcome x trumps outcome y for an agent if the agent receives a higher transfer under outcome x , and all else is equal. We say that a preference ranking \succsim_i is **consistent** with \succeq if $x \succeq y$ implies that $x \succsim_i y$ and $x \succ y$ implies that $x \succ_i y$.

We allow the possibility that different agents have different dominance relations, \succeq_i , and therefore different preference domains. We assume that all rankings in \mathcal{P}_i are consistent with \succeq_i . If $x \succeq_i y$ and $y \succeq_i x$ then x and y are \succeq_i -**equivalent**. Any \succeq_i determines an **equivalence partition** of \mathcal{X} . We refer to each element $[x]_i = \{y \in \mathcal{X} : x \succeq_i y \text{ and } y \succeq_i x\}$ of the equivalence partition as a **payoff** of the agent in question. Consistency implies that each preference ranking in \mathcal{P}_i induces a well-defined preference ranking over payoffs in the natural way: $[x]_i \succsim_i [y]_i$ if $x \succsim_i y$ and $[x]_i \succ_i [y]_i$ if $x \succ_i y$. To avoid unnecessary formalism, we use the same symbol for preferences over payoffs as for preferences over outcomes, and write “payoff x ” for $[x]_i$ and “payoff x obtains” when the realized outcome belongs to $[x]_i$. Unless stated otherwise, we assume in this section that the preference domain \mathcal{P}_i is **rich** in the following sense: the set of induced preferences over payoffs consists of all strict rankings over payoffs.²⁴

²⁴Our use of the term richness shares with other uses of the term in the literature the idea that the domain of preferences contains sufficiently many profiles: if certain preference profiles belong to the domain, then some other profiles belong to it as well (cf. Dasgupta, Hammond, and Maskin (1979) and Pycia (2012)). The more outcome pairs that are comparable by the structural dominance relation \succeq_i , the smaller the resulting preference domain, and hence—in Theorems 5, 7, and 8 below—simple dominance is required for fewer preference types. At one extreme, \succeq_i is an identity relation for each $i \in \mathcal{N}$, agents’ preference domains consist of all strict rankings, and the set of simple mechanisms resembles dictatorships as in Gibbard (1973) and Satterthwaite (1975) and our Corollary 2. At the other extreme, \succeq_i compares all outcomes, each agent is

The framework of rich preference domains is flexible and encompasses many standard economic environments. Examples of rich domains 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. Richness implies that each \mathcal{P}_i consists of all strict preference rankings over \mathcal{X} .
- **Allocating indivisible goods without transfers:** Each $x \in \mathcal{X}$ describes the entire allocation of goods to each of the agents. Each agent has strict preferences over each bundle of goods she may receive, but 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. Richness implies that every strict ranking of these sets belongs to \mathcal{P}_i for each i .

With such examples in mind, we say that an environment is **without transfers** if the dominance relation \succeq_i is symmetric for all i .²⁵ Non-symmetric dominance relations \succeq_i allows us to model transfers: all else equal, having more money dominates having less. Examples of rich domains with transfers include:

- **Social choice with transfers:** Let $\mathcal{X} = \mathcal{Y} \times \mathcal{W}^{\mathcal{N}}$, where \mathcal{Y} is a set of substantive outcomes and $\mathcal{W} \not\subseteq \mathbb{R}$ a (finite) set of possible transfers. 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. The structural dominance relation is then $(y, w) \succeq_i (y', w')$ if and only if $y = y'$ and $w_i \geq w'_i$ (where $w \equiv (w_i)_{i \in \mathcal{N}}$ is the profile of transfers).
- **Auctions:** Let $\mathcal{X} \subseteq \mathcal{N}^{\mathcal{O}} \times \mathcal{W}^{\mathcal{N}}$ where \mathcal{O} is a finite set of goods and $\mathcal{W} \not\subseteq \mathbb{R}$ is a finite set of transfers. Each agent i prefers to win more goods and to pay less rather than more. Denoting by O_i the set of goods allocated to i and writing $O = (O_i)_{i \in \mathcal{N}}$, the structural dominance relation is given by $(O; w) \succeq_i (O'; w')$ if and only if $O_i \supseteq O'_i$ and $w_i \geq w'_i$.

indifferent among all outcomes, and all extensive-form games, with any strategies, are simple. In between these extremes, we have other classes of simple mechanisms, as we explore in this section. We would like to thank referees for these clarifications.

²⁵A 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.

These are just a few examples of settings that fit into our general model. While richness is a flexible assumption, not all preference domains are rich. For instance, domains of single-peaked preferences are typically not rich and Arribillaga, Massó, and Neme (2017) show that our millipede construction does not extend to single-peaked preference domains.

4.1 Obvious Dominance

Obvious dominance (defined above) is the simplicity standard introduced by Li (2017b). Recall that in analyzing obvious dominance we do not need to distinguish between strategies and strategic plans; thus, for simplicity of exposition, we focus on strategies. If a game Γ admits a profile of obviously dominant strategies, then the game and the resulting mechanism (game and strategy profile) are said to be **obviously strategy-proof (OSP)**.

Li’s (2017b) shows that in binary allocation environments with transfers, every OSP mechanism is equivalent to a personal clock auction. In this section, we focus on environments without transfers and show that any OSP game is equivalent to what we call a *millipede game*. Roughly speaking, a millipede game is a clinch-or-pass game similar to a centipede game (Rosenthal, 1981), but in general with more players and more actions (i.e., “legs”) at each node. A simple example of a millipede game in an object allocation environment is a deterministic **serial dictatorship** in which there are no passing moves and all payoffs that are not precluded by the earlier choices of other agents are clinchable (cf. Sections 4.3 and 5).

As a preliminary step to define millipede games, we introduce the following definitions, which apply to any game Γ . Given some history h , we say that payoff x is **possible** for agent i at h if there is a terminal history $\bar{h} \supseteq h$ such that at the outcome associated with \bar{h} , agent i obtains payoff x . We use $P_i(h)$ to denote the set of possible payoffs for i at h . We say that agent i has **clinched** payoff x at history h if at all terminal histories $\bar{h} \supseteq h$, agent i receives payoff x . If i moves at h , takes action $a \in A(h)$, and has clinched x at the history (h, a) , then we call action a a **clinching action**; any action at h that is not a clinching action is called a **passing action**. We denote by $C_i(h)$ the set of all payoffs x that are **clinchant** for i at h ; that is, $C_i(h)$ is the set of payoffs for which there is an action $a \in A(h)$ such that i has clinched x at the history (h, a) . At a terminal history \bar{h} , no agent is called to move and there are no actions; however, it will be notationally useful to define $C_i(\bar{h}) = \{x\}$, where x is the payoff that i 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 while the definition of $C_i(h)$ presumes that i moves at h or h is terminal, the payoff sets $P_i(h)$, $C_i^c(h)$ and $C_i^c(h)$ are well-defined for any h , whether i moves at h or not, and whether h is terminal or not. Finally, consider a history h such that $i_{h'} = i$ for some $h' \not\subseteq h$ 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 $h' \not\subseteq h$ such that $i_{h'} = i$, but $x \notin P_i(h)$. We say payoff x is **previously unclinable** at h if $x \notin C_i^c(h)$.

Given a mechanism $(\Gamma, S_{\mathcal{N}})$ and a type \succ_i , we call strategy $S_i(\succ_i)$ a **greedy strategy** if at any history $h \in \mathcal{H}_i$ it satisfies the following: if the \succ_i -best still-possible payoff in $P_i(h)$ is clinchable at h , then $S_i(\succ_i)(h)$ clinches this payoff; otherwise, $S_i(\succ_i)(h)$ is a passing action. A greedy strategic plan is defined in the same way.²⁶

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 history at which an agent moves, all but at most one action are clinching actions, and following any clinching action, the agent does not move again.
3. At all h , if there exists a previously unclinable payoff x that becomes impossible for agent i_h at h , then $C_{i_h}^c(h) \subseteq C_{i_h}(h)$.

We refer to millipede games with greedy strategies as millipede mechanisms. In a millipede game, it is obviously dominant for an agent to clinch the best possible payoff at h whenever it is clinchable. The last condition of the millipede definition ensures that passing at h is obviously dominant when an agent's best possible payoff at h is not clinchable.

Theorem 5. (Millipedes). *Consider an environment without transfers. Every OSP mechanism is equivalent to a millipede mechanism. Every millipede mechanism is OSP.*

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 (set of) object(s) and leaving the remaining objects to be distributed amongst the remaining agents. Theorem 5 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, we have the following:

²⁶A stronger concept of a greedy strategy would additionally require that when passing, the agent takes an action a such that they are indifferent between the best possible payoffs at h and (h, a) . (Such an action a exists because $P_i(h) = \cup_{a \in A(h)} P_i((h, a))$.) This distinction is immaterial for millipede games, since they have at most one passing action at each history, and all of our results are valid for both concepts of greediness.




Figure 1: An example of a non-dictatorial millipede game in a voting environment with two outcomes, $\mathcal{X} = \{x, y\}$. The obviously dominant (greedy) strategy profile is for any agent to clinch if she is offered to clinch her preferred option among $\{x, y\}$, and otherwise pass.

Corollary 2. *Let each agent’s preference domain \mathcal{P}_i be the space of all strict rankings over outcomes \mathcal{X} . Then, every OSP game is equivalent to a game in which either:*

- (i) *the first agent to move can clinch any possible outcome and has no passing action; or*
- (ii) *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.*

The former case of Corollary 2 is the standard dictatorship, with a possibly restricted set of outcomes. The latter case is a generalization that allows an agent to enforce one of the two outcomes, but not the other, at her turn; see Figure 1 for an example. In particular, this corollary gives an analogue of the Gibbard-Satterthwaite dictatorship result, with no efficiency assumption.

The full proof of Theorem 5 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. We construct this millipede game via the following transformations. Starting with any arbitrary game, we begin by breaking information sets; this only shrinks the set of possible outcomes any time an agent is called to play, which preserves the min/max obvious dominance inequality. For similar reasons, we can shift all of Nature’s moves to the beginning of the game, and so now have a perfect information game Γ' in which Nature moves once, as the first mover.²⁷ Second, if there are two passing actions a and a' at some on-path 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' . Then, we can construct an equivalent game via a transformation in which we add an action that allows i to clinch this payoff already at h by making all such “future choices”

²⁷The first part of this transformation is a special case of Theorem 4. That 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, Lemma A.1 provides the details.




Figure 2: An example of a millipede game with two agents $\{i, j\}$ and 100 objects $\{o_1, o_2, \dots, o_{100}\}$. If the first clinching is in an agent's first 50 moves, then the other agent is given the choice of clinching any object he or she could have clinched previously; if the first clinching is after the clinching agent's first 50 moves, then the other agent is given the choice of clinching any still-available object.

today. We also rely on Li's pruning, in which the actions no type chooses are removed from the game tree, cf. Appendix A.1 and Lemma A.3. We repeat these transformations until there is at most one passing action remaining. The final step of the proof is to show that these transformations give us a millipede game. This last step relies on richness and shows that 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.

4.2 One-Step Dominance

In this section, we analyze the stronger simplicity concept of one-step simple dominance. To see why strengthening of obvious dominance might be useful, recall that obviously dominant strategies may not be intuitively simple; an already discussed stark example is White's winning strategy in chess. As another example, consider a no-transfer object allocation environment and the two-player millipede game in Figure 2. At the first move, type $o_{100} >_i o_1 >_i o_2 >_i \dots >_i o_{99}$ is offered her second-favorite object, o_1 , while her top choice, o_{100} , is possible. The obviously dominant greedy strategy of this type is to pass; however, if she does so, she may not be offered the opportunity to clinch her top object, o_{100} , or even go back to her second-best object, o_1 , until far into the future. Thus, while passing is obviously dominant, comprehending this requires the ability to reason far into the future of the game and to perform lengthy backwards induction.²⁸

The more demanding concept of one-step simplicity eliminates the intuitively complex, yet still formally obviously dominant, strategies such as White winning strategy in chess and

²⁸The first 100 moves of this millipede cannot be substantially shortened because, given the players' greedy strategies, for $k = 1, \dots, 50$, i can obtain o_{k+1} if and only if j 's top choice is $o_{100-k+1}$ or a lower-indexed object, and j can obtain $o_{100-k+1}$ if and only if i 's top choice is o_k or a higher-indexed object.

the greedy strategy in the millipede of Figure 2, while still classifying greedy strategies in serial dictatorships and ascending auctions as simple.

Binary allocation with transfers

Li (2017b) illustrates the usefulness of obvious dominance in the setting of binary allocation with transfers, defined as follows. The set of outcomes is $\mathcal{X} = Y \times \mathbb{R}^{\mathcal{N}}$, where $Y \subseteq \{0, 1\}^{\mathcal{N}}$ is a set of feasible allocations and $\mathbb{R}^{\mathcal{N}}$ is the set of profiles of transfers, one for each agent; a generic allocation is denoted y and a generic profile of transfers $w = (w_i)_{i \in \mathcal{N}}$. In this section, we denote types by $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$, where $0 \leq \underline{\theta}_i < \bar{\theta}_i < \infty$, and assume each agent has preferences represented by a quasilinear utility function: $u_i(\theta_i, y, w) = \theta_i y_i + w_i$.²⁹ This framework captures many important environments of economic interest, including single-unit auctions, procurement auctions, and binary public goods games.

For these environments, Li introduces the class of *personal clock auctions*, which generalize the ascending auction in several ways: agents may face different individualized prices (“clocks”); at any point, there may be multiple quitting actions that allow agents to drop out of the auction, or multiple continuing actions that allow them to stay in the auction; and when an agent quits, her transfer need not be zero. The key restrictions are that each agent’s clock must be monotonic, and that whenever the personal price an agent faces strictly changes, she must be offered an opportunity to quit. The formal definition of a personal clock auction can be found in Appendix A.5, where we also prove Theorem 6.

Li (2017b) shows that in binary allocation settings, OSP games are equivalent to personal clock auctions. Using our new conception of simplicity, we can strengthen this result to show that personal clock auctions are also OSS, and so, perhaps surprisingly, there is no loss in imposing one-step dominance: any OSP-implementable social choice rule is also implementable in one-step dominant strategic collections.

Theorem 6. (*OSS and Personal Clock Auctions*). *In binary allocation settings with transfers, every one-step simple mechanism is equivalent to a personal clock auction with one-step dominant strategic collections. Furthermore, every personal clock auction is one-step simple.*

Because our Corollary 1 shows that any OSS mechanism is also OSP, the first part of the theorem follows from Li’s (2017b) result that any OSP mechanism is equivalent to a personal clock auction with greedy strategies, provided we can find a profile of one-step

²⁹We allow for a continuum of types and transfers here in order to reproduce the binary allocation environment of Li (2017b). Our simplicity concepts extend to this environment when we substitute inf for min and sup for max in our definitions. Richness plays no role in the binary allocation results.

dominant strategic collections that replicates the play of Li’s greedy strategies. We construct these collections in the proof of the second part of the theorem. For the special case of a standard ascending auction for a single good, this can be easily done as follows: at 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. This strategic collection leads to the same outcome as the greedy strategy of staying in at prices weakly below v_i and dropping out at prices strictly above. The collection is one-step dominant because at any p at which the agent’s stay In, she can plan to quit at the next step and get a payoff of 0, which is no worse than quitting now. For personal clock auctions more generally, a key feature is whenever an agent’s price changes, she must be offered an opportunity to quit. This feature allows us to construct one-step dominant strategic plans analogous to those just defined for ascending auctions. The complete argument can be found in the appendix.

Environments without transfers

In environments without transfers, we have seen millipede games that, while OSP, may still be quite complex and require lengthy foresight on the part of the agents (e.g., Figure 2). Imposing the stronger standard of OSS eliminates these complex millipede games, and leaves only games that are monotonic in the following sense: a millipede game Γ is **monotonic** if, for any agent i and any histories h, h' such that: $(h, a^*) \subseteq h'$ where a^* is a passing action at h , $i_h = i$, $i_{h'} = i$ or h' is terminal, and $i_{h''} \neq i$ for any h'' such that $h \not\subseteq h'' \not\subseteq h'$, 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 if an agent passes at h , the next time she moves, she is offered to clinch either (i) everything she could have clinched at h or (ii) everything that was possible, but not clinchable at h . Some millipede games, such as serial dictatorships in which each agent only moves once and has no passing action, are trivially monotonic; for a less trivial example of a monotonic millipede game, see Figure 3 in Section 5. We say that a mechanism is monotonic when the underlying game is.

Theorem 7. (*Monotonic Millipedes*). *In environments without transfers, every one-step simple millipede mechanism is equivalent to a monotonic millipede mechanism. Furthermore, every monotonic millipede mechanism is one-step simple.*

At any history h in a monotonic millipede game, the one-step dominant strategic plan is as follows: if the agent can clinch her top still-possible outcome at h , then she does so; otherwise, the agent passes at h , and for any next-history h' , the strategic plan is to clinch her top possible object in $C_i(h')$. If clause (i) of monotonicity holds, then this is at least as good as anything she could clinch at h (since the clinchable set weakly expands); if clause

(ii) of monotonicity holds, then she obtains her best possible payoff in $P_i(h)$, which is again at least as good as anything that was clinchable at h .

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, 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 may be strictly worse for many moves in the future, before eventually being re-offered what she was able to clinch in the past (or something better). If agents are unable to plan far ahead in the game tree, it may be difficult to recognize that passing is obviously dominant in such a game; 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.

Further, from a practical implementation perspective, monotonic games are also particularly simple for a designer to run dynamically: at each step, the designer only need tell an agent her possible clinching options today, plus that if she passes, at her very next move, her clinchable set will either weakly expand, or she will be offered everything possible that she was not offered today. Such a partial, one-step-at-a-time description is simpler than trying to describe all of the possibilities many moves in the future that would be necessary to implement more complex, non-monotonic OSP games.

4.3 Strong Obvious Dominance, Choice Mechanisms, and Posted Prices

In light of Theorem 1, the strongest simplicity standard in our class is strong obvious dominance. If a game Γ admits a profile of strongly obviously dominant strategic collections, we say that it is **strongly obviously strategy-proof (SOSP)**. Random Priority is SOSP, but ascending auctions are not. Thus, SOSP mechanisms further delineate the class of games that are simple to play, by eliminating millipede games that require even one-step forward-looking behavior. As for obvious dominance, there is a one-to-one correspondence between strongly obviously dominant strategic collections and strongly obviously dominant strategies, and so for simplicity of exposition, in this section we focus on strategies.

Strongly obvious strategy-proof games are particularly simple to play. As SOSP dominance is established 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 and, 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.³⁰ 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. In the next definition and in the proof of this lemma we say that a payoff x is **undominated** (or untrumped) in a subset of payoffs for agent i if there is no payoff y in this subset such that $y \triangleright_i x$.

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 choice 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. When exactly two undominated payoffs are possible for the agent who moves, the agent can be faced with either (i) a set of clinching actions that allow the agent to clinch either of the two payoffs, (ii) a passing action and a set of clinching actions that allow the agent to clinch exactly one of these payoffs. Note that we allow potentially many ways of clinching the same payoff; we can conceptualize the many ways of clinching a fixed payoff as clinching it and sending a message from a predetermined set of messages. Note also that (ii) does not allow the agent to clinch the other payoff.

³⁰The on-path restriction is not needed if we consider the class of “pruned” games in the sense of Li (2017b); cf. Appendix A.1 for how pruning works with OSP, SOSP, and other simplicity concepts.

Theorem 8. (*Sequential Choice*). *Every strongly obviously strategy-proof mechanism is equivalent to a sequential choice mechanism with greedy strategies. Every sequential choice mechanism with greedy strategies is strongly obviously strategy-proof.*

Theorem 8 applies to any rich preference environment, including both those with and without transfers. In an object allocation model without transfers, every SOSP mechanism resembles a sequential dictatorship, in which agents are called sequentially and offered to clinch any object that still can be clinched given earlier clinching choices; they pick their most preferred object and leave the game. The key difference between a sequential choice game and a sequential dictatorship is that at an agent’s turn, she need not be offered all still-available objects.

In environments with transfers, sequential choice games can be interpreted as sequential posted-price games. 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.

Price mechanisms are ubiquitous in practice. Even 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. Posted prices have also garnered significant attention in the computer science community. For instance, computing the optimal allocation in a combinatorial Vickrey auction can be complex even from a computational perspective, and several papers have shown good performance using sequential posted price mechanisms (e.g., Chawla, Hartline, Malec, and Sivan (2010) and Feldman, Gravin, and Lucier (2014)). By formalizing a strategic simplicity-based explanation for the popularity of these mechanisms, our Theorem 8 complements this literature.³¹

5 Random Priority

As an application 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’

³¹Prior economic studies on the focal role of posted prices in mechanism design—e.g., Hagerty and Rogerson (1987) and Copic and Ponsati (2016)—focused on bilateral trade, while our analysis is applicable to any economic environment satisfying our richness assumption.

choices. Random Priority succeeds on three important design dimensions: it is simple to play, efficient, and fair.³² However, this is only a partial explanation of its success, as to now, it has remained unknown whether there exist other such mechanisms, and, if so, what explains the relative popularity of RP over these alternatives (cf. footnote 10). Theorem 9 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.³³ 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.³⁴ 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 9. (*Random Priority*). *An obviously strategy-proof mechanism is symmetric and Pareto efficient if and only if it is equivalent to Random Priority.*

As discussed above, it is well-known Random Priority satisfies OSP, symmetry, and Pareto efficiency. The converse implication is new. Theorem 9 remains true if we replace OSP with OSS, SOSP, or any other of our simplicity standards; this is implied by combining Theorem 9 with Theorems 1 and 8.

³²Pareto efficiency and fairness of RP have been recognized at least since Abdulkadiroğlu and Sönmez (1998), 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.

³³Because 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).

³⁴We 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 first step in proving Theorem 9 is to recognize that it is sufficient to prove it for any uniform randomization over Pareto efficient deterministic millipedes. The reduction to symmetric randomizations follows because every symmetric mechanism is a lottery over symmetric randomizations (details in appendix), and if each of these randomizations is equivalent to Random Priority then so is the lottery over them. The further reduction to randomizations over Pareto efficient millipedes follows from our Theorem 5.

At the core of the remainder of the proof is the 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. The existence of such a bijection implies that uniform randomizations over permutations of a deterministic Pareto-efficient millipede give the same resulting outcome distribution as Random Priority.³⁵ The full construction is lengthy and involved, and its details can be found in the appendix. Here, we provide a simple three-agent example to showcase the general idea.³⁶

Consider the millipede game presented in Figure 3. The game allocates three objects A, B , and C to three agents (or players) 1, 2 and 3. Agent 1 moves first and can clinch one of the objects A and B or pass. The second move is made by agent 2, who either clinches an object (in which case the allocation is fully determined) or passes (the passing move is only possible following a pass by 1). Agent 3 only moves following two passes; this player can then clinch any object. If Agent 3 clinches A or B then the allocation is determined, and if Agent 3 clinches C then Agent 1 can choose between A and B . This game is a millipede (and OSP), and is Pareto efficient for any preference profile.

We can apply any permutation of agents, $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, to permute an entire mechanism. In a serial dictatorship, this corresponds to changing the order in which the agents select. Similarly, in any arbitrary base game each permutation σ creates a **permuted game** in which agent i is given the moves and payoffs of agent $\sigma(i)$. For instance, if in game Γ in such Figure 3 the agents are permuted by σ such that $\sigma(1) = 1$, $\sigma(2) = 3$, and $\sigma(3) = 2$, then the first agent to move is still agent 1 but the second agent to move is agent 3, and agent 2 moves only after two passes of agents 1 and 3.

³⁵The bijection idea was first employed by Abdulkadiroğlu and Sönmez (1998), and has since been used by others (e.g., Pathak and Sethuraman (2011) and Carroll (2014)). Our construction is fairly different from the bijections in the earlier literature, and relies on the properties of millipede games established by us, and on the properties of Pareto efficient OSP mechanisms subsequently obtained by Bade and Gonczarowski (2017).

³⁶For $|\mathcal{N}| = 1$, the equivalence follows from Pareto efficiency. For $|\mathcal{N}| = 2$, the equivalence is implied by Pareto efficiency when agents rank objects differently and it is implied by symmetry when they rank objects in the same way. Cf. Bogomolnaia and Moulin (2001) who also analyze the three-agent case; their approach is not applicable beyond three agents.




Figure 3: An OSP and Pareto efficient game Γ among three players 1, 2 and 3. Clinching move marked by the object being clinched ($A, B,$ or C); passing moves marked “pass.”

For the sake of illustration, suppose that the preferences are such that $A \succ_1 B \succ_1 C$ for agent 1, $A \succ_2 B \succ_2 C$ for agent 2, and $C \succ_3 B \succ_3 A$ for agent 3. We assume that all agents play greedy strategies. Then, under each permutation σ such that agents 1 or 2 are the first movers, game Γ is played as serial dictatorship would be played: the first mover picks their favorite object, A , and the second mover picks their favorite still-available object, thus also determining the allocation of the third mover. In constructing the bijection between permutations of Γ and permutations of serial dictatorships, we map each of the above permutations $\sigma(\Gamma)$ to the corresponding serial dictatorships. As we prove in the appendix, whenever the game starts with several agents choosing clinching moves, we can map it into a serial dictatorship that starts with the same agents moving in the same order.

The mapping of games that involve passing is more subtle. In our example, passing is on the game path if $\sigma(3) = 1$. There are two such permutations: if $\sigma(2) = 2$ then the resulting outcome is $\{(1, A), (2, B), (3, C)\}$, and if $\sigma(2) = 1$ then the resulting outcome is $\{(1, B), (2, A), (3, C)\}$. To what serial dictatorships should we map these two permutations? The unique mapping achieving the bijection—given how we mapped other permutations—maps the first of the two permutations into the serial dictatorship with agents ordered 3, 1, 2 and the second one into the serial dictatorship with agents ordered 3, 2, 1. There is no simple rule of thumb in mapping permutations that entails passing on the path of play—notice e.g. that in the present example the serial dictatorship order is not the order in which the agents move, nor is it the order in which the agents clinch—and the bulk of the proof is devoted to the general construction of such a mapping.

6 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 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 OSS, and that SOSP mechanisms are equivalent to price mechanisms, which are ubiquitous in practice. 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. Along the way, we provide a logically consistent—though limited to simple games—approach to the analysis of agents with limited foresight.

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 and revenues. In environments with transfers, Vickrey (1961), Riley and Samuelson (1981), Myerson (1981), Manelli and Vincent (2010), and Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) show that the efficiency and revenues achieved with Bayesian implementation can be replicated in dominant strategies; thus the accompanying increase in simplicity may come without efficiency and revenue costs. Li (2017b) and our paper advance this insight further and establish that obviously strategy-proof and one-step simple mechanisms can also implement efficient outcomes (and revenue-maximizing 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 and revenue maximization.³⁷ In environments without transfers, however, 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

³⁷For 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 Theorem 8. The same argument in the same example shows that SOSP mechanisms raise less revenue than optimal auctions. 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.

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 Theorem 2 shows that in general, in environments both with and without transfers, imposing more restrictive simplicity standards on the mechanisms limits the set of implementable social choice functions.³⁸

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, while strongly obviously strategy-proof mechanisms are easier still and nearly all participants play them as expected (see Bo and Hakimov, 2020),³⁹ our general theory of simplicity opens new avenues for experimental investigations. For instance, we 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; or as the highest k that still allows the agent to play k -simple strategies correctly. We may similarly define the measure of sophistication of experimental subjects as the highest k that allows the subjects to play k -simple strategies correctly.

In sum, the sophistication of agents may vary across applications, and so it is important to have a range of simplicity standards. For sophisticated agents, a weaker simplicity standard ensures they play the intended strategies, allowing the designer more flexibility on other objectives; however, for less sophisticated agents, a stronger standard of simplicity may need to be imposed to ensure the intended strategies are played, with 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.

³⁸A different approach to the trade-off between simplicity and flexibility was proposed by Li and Dworzak (2020), who study strategy-proofness, obvious strategy-proofness, and strong obvious strategy-proofness. While we evaluate this tradeoff for designers who never confuse the mechanism participants, they evaluate it for designers who can confuse participants. See also work in progress by Catonini and Xue (2021), who study a weakening of one-step simplicity.

³⁹For a test of the first claim see also Breitmoser and Schweighofer-Kodritsch (2019).

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.
- ARORA, S. AND B. BARAK (2009): *Computational Complexity. A Modern Approach*, Cambridge University Press.
- 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,” .
- BERGEMANN, D. AND S. MORRIS (2005): “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- BO, I. AND R. HAKIMOV (2020): “Pick-an-Object Mechanisms,” SSRN 3560372 Working Paper.
- BOGOMOLNAIA, A. AND H. MOULIN (2001): “A New Solution to the Random Assignment Problem,” *Journal of Economic Theory*, 100, 295–328.
- BÖRGERS, T. AND J. LI (2019): “Strategically Simple Mechanisms,” *Econometrica*, 87, 2003–2035.
- 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.

- CATONINI, E. AND J. XUE (2021): “Local Dominance,” Work in Progress.
- 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.
- 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.
- GABAIX, X., D. LAIBSON, G. MOLOCHE, AND S. WEINBERG (2006): “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96, 1043–1068.
- 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.
- KARNI, E. AND M.-L. VIERO (2013): ““Reverse Bayesianism”: A Choice-Based Theory of Growing Awareness,” *American Economic Review*, 103, 2790–2810.
- KE, S. (2019): “Boundedly rational backward induction,” *Theoretical Economics*, 14, 103–134.
- LAIBSON, D. (1997): “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 112, 443–478.

- LI, J. AND P. DWORCZAK (2020): “Are Simple Mechanisms Optimal when Agents are Unsophisticated?” Working Paper.
- 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 (2020): “Asymptotically Optimal Prior-Free Clock Auctions,” *Journal of Economic Theory*, 187, 1–34.
- MACKENZIE, A. (2017): “A Revelation Principle for Obviously Strategy-proof Implementation,” *working paper*.
- MANDAL, P. AND S. ROY (2020): “Obviously Strategy-proof Implementation of Assignment Rules: A New Characterization,” .
- MANELLI, A. M. AND D. R. VINCENT (2010): “Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model,” *Econometrica*, 78, 1905–1938.
- MASKIN, E. AND J. TIROLE (1999): “Unforeseen Contingencies and Incomplete Contracts,” *The Review of Economic Studies*, 66, 83–114.
- 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.
- PATHAK, P. A. AND T. SÖNMEZ (2008): “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *American Economic Review*, 98, 1636–1652.
- PEARCE, D. G. (1984): “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52, 1029–1050.

- 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 (2017): “Incentive Compatible Allocation and Exchange of Discrete Resources,” *Theoretical Economics*, 12, 287–329.
- (2020): “Arrovian Efficiency and Auditability in the Allocation of Discrete Resources,” CEPR Discussion Paper DP15377.
- 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.
- SCHWARTZSTEIN, J. AND A. SUNDERAM (2021): “Using Models to Persuade,” *American Economic Review*, 111, 276–323.
- SHAPLEY, L. AND H. SCARF (1974): “On Cores and Indivisibility,” *Journal of Mathematical Economics*, 1, 23–37.
- SHIMOJI, M. AND J. WATSON (1998): “Conditional Dominance, Rationalizability, and Game Forms,” *Journal of Economic Theory*, 83, 161–195.

- 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.
- SPENNER, P. AND K. FREEMAN (2012): “To Keep Your Customers, Keep It Simple,” *Harvard Business Review*, May.
- TROYAN, P. (2019): “Obviously Strategy-Proof Implementation of Top Trading Cycles,” *International Economic Review*, 60.
- TROYAN, P., D. DELACRETAZ, AND A. KLOOSTERMAN (2020): “Essentially Stable Matchings,” *Games and Economic Behavior*, 120, 370–390.
- VICKREY, W. (1961): “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37.
- WILSON, R. (1987): “Game-Theoretic Approaches to Trading Processes,” in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, Cambridge: Cambridge University Press, 33–77.
- 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,” .

A Supplementary Appendix: Proofs (For Online Publication)

A.1 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.

In general, for all simple dominance concepts we study, we say that a game is pruned if, for any agent i and any information set $I \in \mathcal{I}_i$, either there are no precedent information sets in \mathcal{I}_i or there is a type of agent i and information set $I' \in \mathcal{I}_i$ such that I' is on the path of play for this type and I is simple from the perspective of I' .

A.2 Proof of Theorem 2

In light of Corollary 1, it is sufficient to prove the result for $k < \infty$ and $k' = k + 1$. We first construct a social choice rule that cannot be one-step implemented, but can be k -step implemented for any $k \geq 2$. We then provide the details how to extend the argument to show that there exist $(k + 1)$ -step implementable rules that are not k -step implementable for any $k \geq 1$. The case of $k = 0$ is dealt with separately at the end.

Consider an object allocation environment without transfers in which agents demand exactly one object each. There are at least three agents i, j, ℓ and the objects included in the game Γ are shown in Figure 4. Each branch of the game tree represents a clinching action where the agent clinches the labeled object $(x, \tilde{x}, \text{etc.})$. The notation such as “ $\ell \rightarrow \gamma$ ” below terminal nodes denotes that agent ℓ is assigned to object γ at this node, without needing to take any action. Note that in this game, if i clinches x at the first move, then the game immediately ends with j assigned α_j and ℓ assigned α_ℓ , and further, this is the only terminal history at which j receives α_j and ℓ receives α_ℓ . Similarly, there are objects β_ℓ, γ_ℓ , and δ_ℓ that agent ℓ receives only at the denoted terminal histories, and nowhere else in the game.

We claim that the rule implemented by the mechanism $(\Gamma, S_{\mathcal{N}})$, where $S_{\mathcal{N}}$ is a profile of greedy strategies, is k -step implementable for any $k \geq 2$, but cannot be OSS-implemented. It is straightforward to check that Γ is k -step implementable for any $k \geq 2$; note that in particular, this implies that Γ is OSP. It is also easy to check that Γ itself is not implementable




Figure 4: A millipede game that cannot be OSS-implemented.

for $k = 1$, or in other words, Γ is not OSS: at i 's second move on the passing path, OSS is violated and so at i 's first move, the type of i that ranks $w \succ_i x \succ_i z$ has no one-step simple strategic plan. In order to show that there is no OSS mechanism that is equivalent to Γ (i.e., implements the same social choice rule), suppose, by way of contradiction, that there is such mechanism with game $\tilde{\Gamma}$ and a profile of OSS strategic plans. Let \tilde{S} be the profile of strategies in $\tilde{\Gamma}$ induced by the strategic plans; by Theorem 1, this profile is obviously dominant.

The proof proceeds in a series of steps, which we label 1.1-1.6. (The labeling $k.1 - k.6$ is used because, after proving the result for $k = 1$, we will use analogues of these steps to complete the proof for arbitrary k .)

Step 1.1. In $\tilde{\Gamma}$, the first mover must be i , and x must be guaranteeable for i . Furthermore, at the empty history, w and z are not guaranteeable for i , but there is a unique action after which w and z are possible. This action is taken by all types of player i that rank either w or z first; we call this action i 's focal action.

Proof of Step 1.1. First notice that i must be the first mover. Indeed, in mechanism Γ , agent j receives α_j if and only if agent i prefers x to w and z . Assume that, under $\tilde{\Gamma}$, agent j moves first. Something must be guaranteeable for agent j at this history, say λ .⁴⁰ If $\lambda = \alpha_j$, then we have non-equivalence when j prefers α_j the most and agent i does not prefer x to w and z . If $\lambda \neq \alpha_j$, then, we have non-equivalence when j prefers λ the most and i prefers x to w and z . Therefore, the first mover cannot be j . As the same argument works for agent ℓ , the first mover must be i .

Second, note that equivalence implies that i obtains x for any preference profile such that i prefers x the most, and therefore, x is guaranteeable at the first move in $\tilde{\Gamma}$. Analogously, w

⁴⁰That something must be guaranteeable follows because each history has at least two actions, and in any OSP game, there can be at most one action such that there is some payoff that is possible, but not guaranteeable (see the proof of Theorem 5).

and z must be possible but not guaranteeable for i at the first move. To see that w cannot be guaranteeable, note that if it were, i would receive w for all preference profiles where she ranked it first, which is not the case in Γ , and so equivalence is violated; the same holds for z . By equivalence, both w and z are possible for i , i.e., $w, z \in P_i(h)$. Further, there must be a unique action a^* such that $w, z \in P_i((h, a^*))$. If there were two actions a_1, a_2 such that w were possible after both, then any type that prefers w the most would have no obviously dominant action, since w is not guaranteeable; the same holds for z . Therefore, each of w and z are possible after exactly one action, label them a_w and a_z . If $a_w \neq a_z$, then any type that ranks w first and z second would have no obviously dominant action.⁴¹ Therefore, $a_w = a_z$; we call this action i 's focal action. Since w and z are possible following only the focal action, all types that rank either w or z first must select it. This completes the proof of Step 1.1.

Step 1.2. In $\tilde{\Gamma}$, at the history following the first focal action by i , agent j moves. At this history, both \tilde{x} and x are guaranteeable for j , while a is not guaranteeable. Further, there is a unique action after which a is possible, and this action is taken by all types of j who rank a first; we call this action j 's focal action.

Proof of Step 1.2. Since, per Step 1.1, both w and z are possible for i following the focal action, the focal action cannot lead to a terminal history, and so there must be an agent who moves. We start by showing that the mover must be j . Note that in Γ , agent ℓ receives β_ℓ if and only if agent i prefers either w or z to x , and agent j prefers \tilde{x} the most out of $\{x, \tilde{x}, a\}$. Suppose that i prefers either w or z to x , so that i follows the focal action at the initial history. By the same logic as in Step 1.1, if agent ℓ is the next mover, she must be able to guarantee some payoff, say γ . If $\gamma = \beta_\ell$, this would lead to a non-equivalence when ℓ ranks γ first and j ranks x first. If $\gamma \neq \beta_\ell$, then we have a non-equivalence when ℓ ranks γ first and j ranks \tilde{x} first. Therefore, ℓ cannot be the next mover, and neither can be i (as i just moved) and so it must be j .

The equivalence of Γ and $\tilde{\Gamma}$ implies that for any profile such that i prefers w or z over x and j prefers x the most, j receives x . Because, per Step 1.1, all types of i take the focal action in $\tilde{\Gamma}$, we conclude that following i 's focal action, j must be able to guarantee himself x . The same argument applies for \tilde{x} . Similarly, equivalence implies that there must be an action for j such that a is possible. Outcome a cannot be guaranteeable for j , because if it were, then j would receive a for all preference profiles where i ranks w or z first and j ranks a first, which is not the case in Γ . By an argument similar to Step 1.1, there cannot be any

⁴¹Since w is not guaranteeable and z is not possible after a_w , the worst case from any strategy that selects a_w is strictly worse than z , which is possible from a_z . Similarly, since w is not possible following a_z , the worst case is strictly worse than w , which is possible from a_w . Note that an analogous argument would apply to any type that ranks z first and w second.

other actions after which a is possible, and all types of j that rank a first must select this action. We label this action j 's focal action.

Step 1.3. In $\tilde{\Gamma}$, following i 's focal action and j 's focal action, there might be any finite number of consecutive histories at which i and j move. At these histories where i moves, i can clinch x , but neither w nor z are guaranteeable, and there is a unique action (the focal action) after which w and z are possible and that is taken by all types of i that rank w or z first. At these histories where j moves, both \tilde{x} and x are guaranteeable, but a is not guaranteeable, and there is a unique action (the focal action) after which a is possible and is taken by all types of j that rank a first. Following this sequence of focal actions, agent ℓ moves.

Proof of Step 1.3. Since, per Step 1.2, a is possible, but not guaranteeable following j 's focal action, the focal action cannot lead to a terminal history, and so must lead to a history at which an agent moves. As j just moved, the next mover must be either i or ℓ . If the next mover is i , as the history is on-path for all types of i who prefer w or z over x , the OSS property of $\tilde{\Gamma}$ implies that either x or else both w and z are clinchable for i . Equivalence implies that neither w nor z can be clinchable for i : if w were clinchable, then i receives w for all profiles such that i prefers w the most and j prefers a the most, which is not the case in Γ ; an analogous argument applies for z . Therefore, x must be clinchable. Furthermore, w and z are possible but not guaranteeable for i , and so, as in Step 1.1, OSP implies that there is a unique action after which both w and z are possible, and all types that rank either w or z first takes this action (note that these types must have taken the focal action at i 's initial move, and so are on-path); we call this action the focal action.

Following the focal action by i , the next mover must be j or ℓ . If it is j , then an analogous argument as for i shows that this agent must have both x, \tilde{x} clinchable, and that there must be a unique action after which a is possible but not guaranteeable; we call it the focal action.

Following j 's focal action, the next move is by i or ℓ . If it is by i then the above argument applies again. We might then have a sequence of moves by i and j to which the above two arguments apply. As the game is finite and at the end of every focal action in the sequence more than one outcome is possible, the focal path of the game must lead to a history at which ℓ is called to play. This proves Step 1.3.

Step 1.4. In $\tilde{\Gamma}$, at ℓ 's move following the sequence of focal actions described in Step 1.3, both \tilde{a} and a are guaranteeable for ℓ , while neither c nor x are guaranteeable. There is also a unique action (the focal action) after which c and x are possible for ℓ . This action is taken by all types of ℓ that rank c first.

Proof of Step 1.4. Using arguments similar to Step 1.2, equivalence implies that at ℓ 's move, both \tilde{a} and a are guaranteeable for ℓ , while neither c nor x are guaranteeable, but

both c and x are possible following a unique action that is taken by all types of agent ℓ that rank c first. Since c is not guaranteeable, this action cannot lead to a terminal history. Since c is possible following only the focal action, all types of ℓ that rank c first must select this action. This proves Step 1.4.

Step 1.5. In $\tilde{\Gamma}$, following the above sequence of focal actions that ends with the first focal action by ℓ , there might be any finite number of consecutive histories at which j and ℓ move. Each of these histories has a unique action (the focal action) after which a is possible for j 's moves, and c and x are possible for ℓ 's moves. All types of j that rank a first and all types of ℓ that rank c first take their respective focal actions. Following this sequence of focal actions, the next mover is i .

Proof of Step 1.5. Since there are multiple possible outcomes for k following her focal action, the focal action cannot lead to a terminal history. As k just moved, the next mover must be either i or j . First consider the case in which j moves next. The OSS property implies that either both x and \tilde{x} are clinchable for j , or a is clinchable for j . Consider the latter case. If this were true, then under a preference profile where i prefers w most and z second, j prefers a most, and ℓ prefers c most, j would receive a , which is not the case in Γ . Therefore, j must be able to clinch x and \tilde{x} . By equivalence, a must be possible for j , but not guaranteeable, and so once again there must be a unique focal action after which a is possible and that is taken by all types of j that prefer a the most (note that all of these types have passed at j 's prior moves, and so are on-path). Following the focal action, the next mover is i or ℓ . If it is ℓ , then an analogous argument implies that ℓ must be able to clinch a and \tilde{a} , with c possible but not guaranteeable following a unique focal action. There may again be a sequence of moves by j and ℓ for which this argument can be repeated. As the game is finite and at the end of every focal action more than one outcome is possible, the focal path must lead to a history at which i is called to play. This proves step 1.5.

Step 1.6. In $\tilde{\Gamma}$, at i 's move following the sequence of focal actions described in Step 1.5, x is not clinchable for i .⁴² At this move, there is a unique action (the focal action) after which w is possible for i ; the focal action is also the unique action after which x is possible for i . This focal action is taken by all types of i that rank w first.

Proof of Step 1.6. By way of contradiction, suppose x is clinchable for i . Then OSP implies that in the continuation game following i 's clinching of x , there must be a terminal history at which j receives a : if there were not, then the type of j that prefers a the most and x second would have no obviously dominant action at j 's prior moves. At this terminal history, agent ℓ must be assigned something other than x (which was assigned to i) or a (which was assigned to j). But then, the type of ℓ that prefers x the most and a second has

⁴²The argument shows that x not only is not clinchable for i but also not guaranteeable.




Figure 5: A 3-step simple mechanism for which no equivalent 2-step simple mechanism exists.

no obviously dominant action at ℓ 's prior moves, which is a contradiction.⁴³

An analogous argument to that which showed that there is a unique action after which w is possible for i in Step 1.1, tell us that there is a unique action (the focal action) after which w is possible for i . By OSP, types of i ranking w first take this action. An analogous argument shows that the focal action is the unique action after which x is possible.

Finishing the proof for $k = 1$.

As the previous step shows that x is not clinchable at the move of i considered there, OSS implies that both w and z must be clinchable for i . This implies that for preference profiles such that i ranks w first and x second, j ranks a first, and k ranks c first, agent i is assigned w . However, under such profiles in Γ , i receives x , which is a contradiction to equivalence.

Thus far, we have shown that social choice rule implemented by the game Γ can be k -step implemented for any $k \geq 2$, but cannot be OSS-implemented. The argument that for any $k \geq 1$, there exist rules that can be $(k + 1)$ -step implemented, but cannot be k -step implemented, follows the same steps and uses similar counterexample games as the $k = 1$ game in Figure 4. For $k = 2$, our counterexample is the game $\Gamma^{(2)}$ in Figure 5. This game is similar in structure to that of Figure 4, but has the following additions:

(i) In the subgame following i passing and j clinching x at its first move, we add the possibility of i clinching z' . In this way we assure that i can then clinch any possible and not previously clinchable object.⁴⁴

⁴³Note that by equivalence, x must be possible for ℓ at these prior moves, since in Γ , k receives x for type profiles such that i ranks w first, j ranks a first, and ℓ ranks x first.

⁴⁴Note that this property and the property reestablished in (ii) were also true in game Γ in Figure 4 and these two modifications simply reestablish these properties for the game $\Gamma^{(2)}$ in Figure 5 in which z' becomes

(ii) In the subgame following i and j passing and ℓ clinching a at its first move, we add the possibility of i clinching z' (following which j can clinch x and \tilde{x}). In this way we assure that i can clinch any possible and not previously clinchable object.

(iii) Following i 's pass at its second move on the focal path, we add a node at which ℓ can clinch two new objects a' and \tilde{a}' (following the clinching of a' , agent i can clinch any possible not previously clinchable object, and then j can clinch any previously clinchable object; following the clinching of \tilde{a}' , agent j can clinch any previously clinchable object, and then following the clinching of x agent i can clinch any possible but not previously clinchable objects while following the clinching of \tilde{x} agent i can clinch any previously clinchable object).

(iv) Following the pass at the added node for ℓ , we add a node at which i can clinch an additional object z' . Following i clinching z' , ℓ and then j can clinch any object they could clinch previously).

To generate the counterexample for arbitrary k , we recursively create game $\Gamma^{(k)}$ by adding to game $\Gamma^{(k-1)}$ further objects $z^{(k)}$, $a^{(k)}$, and $\tilde{a}^{(k)}$, and then adding the analogues of subgames (i)-(iv). In the analogues of subgames (i)-(ii), we now allow i to additionally clinch $z^{(k)}$; in the analogue of (iii), $a^{(k)}$ and $\tilde{a}^{(k)}$ play the roles of a and \tilde{a} , and in the analogue of (iv) $z^{(k)}$ plays the role of z .

We can then suppose that $\tilde{\Gamma}^{(k)}$ with a profile of strategic plans is a k -step simple mechanism equivalent to $\Gamma^{(k)}$ with greedy strategic plans. The proof then begins by repeating steps 1.1-1.6, with the only change being that $\Gamma^{(k)}$ plays the role of Γ and $\tilde{\Gamma}^{(k)}$ plays the role of $\tilde{\Gamma}$. Then, we continue with the addition of steps $k'.3-k'.6$ for $k' = 2, 3, \dots, k$. Each step $k'.3-k'.6$ is analogous to the corresponding step 1.3-1.6 from above, except that $a^{(k)}$ plays the role of a , $\tilde{a}^{(k)}$ plays the role of \tilde{a} , and $z^{(k)}$ plays the role of z . Finally, the proof for arbitrary k concludes with a final step that is the direct analogue of the finishing step for $k = 1$, except that we apply k -step simplicity instead of OSS.

Last, the above arguments only apply for $k \geq 1$. The result for $k = 0$ follows from Theorems 6 and 8, applied to a single-unit auction with transfers. Theorem 6 shows that in such a setting, personal clock auctions are efficient and OSS, while Theorem 8 implies that an efficient, SOS ($k = 0$) mechanism does not exist when there are at least two agents whose valuations are drawn iid from at least 3 values (see also footnote 37). ■

A.3 Proof of Theorem 3

The proof builds on and simplifies the proof of the similar result for OSP in Li (2017b). For one direction of implication, suppose the strategic plan S_{i,I^*} is simply dominant from

possible for i .

the perspective of $I^* \in \mathcal{I}_i$ in Γ . Then any outcome that is possible after playing S_{i,I^*} at all information nodes $I \in \mathcal{I}_{i,I^*}$ is weakly better than any outcome that is possible after playing $S'_i(I^*) \neq S_{i,I^*}(I^*)$ in Γ , and hence the analogue of this “weakly better” comparison applies to the counterparts of these actions in any game Γ' that is indistinguishable from Γ from the perspective of i at I^* (by condition (3) of indistinguishability). Hence, in any such Γ' , every strategy S'_i that calls for playing the counterparts of actions $S_{i,I^*}(I)$ for counterparts of all $I \in \mathcal{I}_{i,I^*}$ weakly dominates any strategy S''_i that does not call for playing the counterpart of $S_{i,I^*}(I^*)$ at the counterpart of I^* .

For the other direction of implication, fix information set I^* at which i moves, preference ranking \succeq_i of agent i , and a partial strategic plan S_{i,I^*} such that in every game Γ' that is indistinguishable from Γ from the perspective of agent i at I^* , any strategy S'_i that plays counterparts of $S_{i,I^*}(I)$ for all counterparts of $I \in \mathcal{I}_{i,I^*}$ weakly dominates any strategy S''_i that plays at the counterpart of I^* another action than the counterpart of $S_{i,I^*}(I^*)$. Our goal is to show that any outcome that is possible when i follows S_{i,I^*} at information sets \mathcal{I}_{i,I^*} is \succeq_i -weakly preferred to any outcome that is possible after i plays any $a \neq S_{i,I^*}(I^*)$ at I^* in game Γ . To prove it consider Γ' that differs from Γ only in that all moves of agent i and other agents that follow history h^* but are not in \mathcal{I}_{i,h^*} are made by Nature instead of the party making them in Γ and that Nature puts positive probability on all its possible moves. Notice that such Γ' is indistinguishable from Γ from the perspective of i at I^* . As in Γ' any strategy that selects counterparts of S_{i,I^*} at any counterpart of $I \in \mathcal{I}_{i,I^*}$ weakly dominates any strategy S''_i that selects a at the counterpart of I^* , we conclude from condition (3) of indistinguishability that, in Γ , any outcome that is possible after i follows S_{i,I^*} at information sets in \mathcal{I}_{i,I^*} is weakly better than any outcome that is possible following a . ■

A.4 Proof of Theorem 5

We break Theorem 5 into two propositions. We start by proving that millipede games are OSP (Proposition A.1), and then prove that every OSP game is equivalent to a millipede game (Proposition A.2).

Section 4 introduces the notions of possible and clinchable payoffs at a history h , and the sets of such payoffs, denoted $P_i(h)$ and $C_i(h)$, respectively. For the proof, we will also need the notion of a guaranteeable payoff: a payoff x is **guaranteeable** for i at h if there is some continuation strategy S_i such that i receives payoff x at all terminal histories $\bar{h} \supseteq h$ that are consistent with i following S_i . We use $G_i(h)$ to denote the set of payoffs that are guaranteeable for i at history h .

Proposition A.1. *Millipede games with greedy strategies are obviously strategy-proof.*

Proof. Let Γ be a millipede game. For a set X of payoffs of agent i and a type \succ_i , let $Top(\succ_i, X)$ be the best payoff in X according to preferences \succ_i . Consider some profile of greedy strategies $(S_i(\cdot))_{i \in \mathcal{N}}$. If $Top(\succ_i, C_i(h)) = Top(\succ_i, P_i(h))$, then clinching a 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 following $S_i(\succ_i)$ such that $Top(\succ_i, C_i(h)) \neq Top(\succ_i, P_i(h))$, 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 continuing to follow $S_i(\succ_i)$ at all future nodes.

First, note that $x_W(h) \succeq_i x_W(h')$ for all $h' \not\subseteq h$ such that $i_{h'} = i$. Since passing is obviously dominant at all $h' \not\subseteq h$, 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 and all other actions are clinching actions, we have $x_C(h) \succ_i x_W(h)$; further, since $Top(\succ_i, C_i(h)) \neq Top(\succ_i, P_i(h))$, there must be some $x' \in P_i(h) \setminus C_i(h)$ such that $x' \succ_i x_C(h) \succ_i x_W(h)$. The above implies that $x' \succ_i x_C(h) \succ_i x_C(h')$ for all $h' \not\subseteq h$ 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 that is consistent with $S_i(\succ_i)$ and ends in a terminal history \bar{h} at which i receives $x_W(h)$. For every $x_m \in X_0$, let h_m denote the earliest history on this path such that $x_m \notin P_i(h_m)$ and either (i) $i_{h_m} = i$ or (ii) h_m is terminal. Note that because i is ultimately receiving payoff $x_W(h)$, such a history h_m exists for all $x_m \in X_0$. Let \hat{h}_{-m} be the earliest history at which i moves and at which all payoffs strictly preferred to x_m are no longer possible.

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

Proof of claim. 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, $x_1 = Top(\succ_i, P_i(h))$; since $h' \supset h$ implies that $P_i(h') \subseteq P_i(h)$, we have that $x_1 = Top(\succ_i, P_i(h'))$ as well. Since $x_1 \in C_i(h')$ by supposition, greedy strategies direct i to clinch x_1 , which contradicts that she receives $x_W(h)$.⁴⁵

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 $x_{m'} \succ_i x_m$ be a payoff

⁴⁵ If h' is terminal, then, even though i takes no action at h' , according to our notational convention we define $C_i(h') = \{x_1\}$, which also contradicts that she receives payoff $x_W(h)$.

that becomes impossible at $\hat{h}_{-m} \subseteq \bar{h}$; if such payoff does not exist then the argument of the paragraph above applies. There are two cases:

Case (i): $h' \not\subseteq \hat{h}_{-m}$. This is the case in which x_m is clinchable while there is some strictly preferred payoff $x_{m'} \succ_i x_m$ that is still possible. By assumption, all $\{x_1, \dots, x_{m-1}\}$ are previously unclinched at \hat{h}_{-m} , and so $x_{m'}$ is previously unclinched at \hat{h}_{-m} . By definition of a millipede game (part 3), we have $x_m \in C_i(\hat{h}_{-m})$. Thus, x_m is the best remaining payoff at \hat{h}_{-m} , and is clinchable, and so greedy strategies direct i to clinch x_m at \hat{h}_{-m} , which contradicts that she receives $x_W(h)$ (as in footnote 45, the argument still applies if \hat{h}_{-m} is a terminal history).

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 $x_W(h)$. ■

To finish the proof, let $\hat{h} = \max\{h_1, h_2, \dots, h_M\}$ (ordered by \subset); in words, \hat{h} is the earliest history on the path to \bar{h} at which no payoffs in X_0 are possible any longer. Let \hat{x} be a payoff in X_0 that becomes impossible at \hat{h} . The claim shows that no $x \in X_0$ is clinchable at any $h' \subseteq \hat{h}$, and so we can further conclude that \hat{x} is previously unclinched at \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 $x_W(h)$ (as in footnote 45, the argument still applies if \hat{h} is a terminal history). ■

Proposition A.2. *Every obviously strategy-proof mechanism (Γ, S_N) is equivalent to a millipede mechanism.*

Proof. Given Li's pruning principle (see Subsection A.1), Proposition A.2 follows directly from Lemmas A.1, A.3, and A.4 that we state and prove below.⁴⁶

Lemma A.1. *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. That every OSP game is equivalent to an OSP game with perfect information is implied by our more general Theorem 4. To show that we can furthermore assume that Nature moves at most once, as the first mover, consider a perfect-information game Γ . 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

⁴⁶We 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 A.3 below).

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 $A(h)$ 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. \blacksquare

Lemma A.2. *Let Γ be an obviously strategy-proof game of perfect information that is pruned with respect to the obviously dominant strategy profile $(S_i(>_i))_{i \in \mathcal{N}}$. 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 intersection. By pruning, some type $>_i$ is following some strategy such that $S_i(>_i)(h) = a_1^*$ that results in a payoff of x at some terminal history $\bar{h} \supseteq (h, a_1^*)$. Note that $Top(>_i, P_i(h)) \neq x$, 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 $Top(>_i, P_i(h)) = y$. Note that $y \notin G_i(h)$ (or else it would not be obviously dominant for type $>_i$ to play a strategy such that x is a possible payoff). Further, we must have $y \in P_i(h_1^*)$ 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 $>_i$, which contradicts that $S_i(>_i)(h) = a_1^*$; given the former, if $y \in P_i(h_2^*)$, then once again a_1^* would not be obviously dominant for type $>_i$. Now, again by pruning, there must be some type $>'_i$ such that $S_i(>'_i)(h) = a_2^*$ that results in payoff x at some terminal history $\bar{h} \supseteq (h, a_2^*)$. By similar reasoning as previously, $Top(>'_i, P_i(h)) \neq x$, and so $Top(>'_i, P_i(h)) = z$ 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 $z \notin G_i(h)$. It is without loss of generality to consider a type $>'_i$ such that $Top(>'_i, P_i(h) \setminus \{z\}) = y$. 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 (since $y \notin P_i(h_2^*)$ and $z \notin G_i(h)$), 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 $y \notin P_i(h_1^*)$. Next, for

all types of agent i that reach h , it must be that $Top(>_i, P_i(h)) \neq x, y$. To see why, assume there were a type that reaches h such that $Top(>_i, P_i(h)) = x$. Then, by richness, there is a type that reaches h such that $Top(>_i, P_i(h) \setminus \{x\}) = y$. But, note that this type has no obviously dominant action at h . An analogous argument applies switching x with y .

Now, by pruning, there is some type that reaches h that plays a strategy such that $S_i(>_i)(h) = a_1^*$ and x is a possible payoff. Let $Top(>_i, P_i(h)) = z$ for this type, where, as just noted, $z \neq x, y$. The fact that $S_i(>_i)(h) = a_1^*$ implies that $z \in P_i(h_1^*)$ and $z \notin G_i(h)$; if either of these were false, it would not be obviously dominant for this type to play a strategy such that $S_i(>_i)(h) = a_1^*$ and x is a possible payoff. In other words, $z \in PnG(h)$ 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 $>_i$ such that $Top(>_i, P_i(h)) = z$ and $Top(>_i, P_i(h) \setminus \{z\}) = y$. Note that this type has no obviously dominant action at h , which is a contradiction. ■

Clinching actions are those for which i 's payoff is completely determined after following the action. Lemma A.2 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 . This still 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 A.3. *For any OSP game Γ , there exists an equivalent OSP game Γ' 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 .*
- (iii) *If $P_i(h) = G_i(h)$, then all actions in $A(h)$ are clinching actions and $i_{h'} \neq i$ for any $h' \not\supseteq h$.*

Proof. Given an OSP mechanism (Γ, S_N) , begin by using Lemma A.1 to construct an equivalent OSP game of perfect information in which Nature moves only at the initial history (if at all). Further, prune this game according to the obviously dominant strategy profile S_N . With slight abuse of notation, we denote this pruned, perfect information mechanism by (Γ, S_N) . Consider some history h of Γ at which the mover is $i_h = i$. By Lemma A.2, 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 $>_i$ that does not choose a^* guarantees the best possible

outcome in $P_i(h)$ for type \succ_i . Define the set

$$\mathcal{S}_i(h) = \{S_i : S_i(h) \neq a^* \text{ and at all terminal } \bar{h} \text{ consistent with } S_i, i \text{ receives the same payoff}\}.$$

In words, 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 $P_i((h, a^*)) \notin G_i(h)$ 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 outcome 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

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.

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 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 the pruning principle. ■

Lemma A.4. *Let $(\Gamma, S_{\mathcal{N}})$ be an obviously strategy-proof mechanism that satisfies the conclusions of Lemmas A.1 and A.3. At all h , if there exists a previously unclinched payoff z that becomes impossible for agent i_h at h , then $C_{i_h}^c(h) \subseteq C_i(h)$.*

Proof. Let h be any earliest history where some agent i moves such that there is a previously unclinched payoff z that becomes impossible at h (the case for terminal histories will be dealt with separately below). This means that i moves at some strict subhistory $h' \subsetneq h$ and the following are true:

- (a) $z \notin P_i(h)$
- (b) $z \in P_i(h')$ for all $h' \subsetneq h$ such that $i_h = i$
- (c) $z \notin C_i^c(h)$

Points (b) and (c) imply that z is possible at every $h' \subsetneq h$ at which i is called to move, but it is not clinchable at any of them. This means that for any type of agent i that ranks z first, any obviously dominant strategy must have the agent choosing the unique passing action at all $h' \subsetneq h$.

We want to show that $C_i^c(h) \subseteq C_i(h)$. Towards a contradiction, assume that $C_i^c(h) \not\subseteq C_i(h)$, and let $x \in C_i^c(h)$ but $x \notin C_i(h)$. Consider a type \succ_i that ranks z first and x second. By the previous paragraph, this type must be playing some strategy that passes at any $h' \subsetneq h$, and so h is on the path of play for type \succ_i . Since $z \notin P_i(h)$ and $x \notin C_i(h)$, by Lemma A.3, part (ii), the worst case outcome from this strategy is some y that it is strictly worse than both z and x according to \succ_i . However, we also have $x \in C_i(h')$ for some $h' \subsetneq h$, and so

the best case outcome from clinching x at h' 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 with \bar{h}). Recall that for terminal histories, we define $C_i(\bar{h}) = \{y\}$, where y is the payoff that obtains at h for i . Towards a contradiction, assume that there is some $x \in C_i(h')$ for some $h' \not\subseteq \bar{h}$ but $x \notin C_i(\bar{h})$. Note that (i) $z \neq y$, by (a); (ii) $z \neq x$, by (c); and (iii) $x \neq y$, since $x \notin C_i(\bar{h})$. In other words, x, y , and z are all distinct payoffs for i . Thus, consider the type \succ_i that ranks z first, x second, and y third, followed by all other payoffs. By (b) and (c), z is possible at every $h'' \not\subseteq \bar{h}$ at which i moves, but is not clinchable at any such h'' . Thus, any obviously dominant strategy for type \succ_i must have agent i passing at all such histories. This implies that y is possible for this type from this strategy. However, at h' , i could have clinched x , and so the strategy is not obviously dominant, a contradiction.

■

A.5 Proof of Theorem 6

Before proving the theorem, we first present a formal definition of a personal clock auction. Given some (perfect-information) game Γ , define outcome functions g as follows: $g_y(\bar{h}) \subseteq \mathcal{N}$ is the set of agents who are in the allocation \bar{y} that obtains at terminal history \bar{h} (that is, $i \in g_y(\bar{h})$ if and only if $\bar{y}_i = 1$), and $g_{w,i}(\bar{h}) \in \mathbb{R}$ 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.⁴⁸

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

In Transfer Falls: there exists a fixed transfer $\bar{w}_i \in \mathbb{R}$, a going transfer $\tilde{w}_i : \{h_i : h_i^* \subseteq h_i \subseteq \bar{h}\} \rightarrow \mathbb{R}$ 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 $g_{w,i}(\bar{h}) = \inf\{\tilde{w}_i(h_i) : h_i^* \subseteq h_i \not\subseteq \bar{h}\}$.
- If $\bar{h} \not\supseteq (h, a)$ for some $h \in \mathcal{H}_i$ and $a \in A^q$, then $i \notin g_y(\bar{h})$.
- $A^q \cap A(h_i^*) \neq \emptyset$

⁴⁸In light of our notion of equivalent mechanisms and Theorem 4, the presentation below is a slight simplification of Definition 15 of Li (2017b) that removes imperfect information: 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 also valid for the minor correction of the definition Li gives in a corrigendum available from his website. For further details on this latter point, see footnote 49 below.

- For all $h'_i, h''_i \in \{h_i \in \mathcal{H}_i : h_i^* \subseteq h_i\}$:
 - If $h'_i \not\subseteq h''_i$, then $\tilde{w}_i(h'_i) \geq \tilde{w}_i(h''_i)$
 - If $h'_i \not\subseteq h''_i$, $\tilde{w}_i(h'_i) > \tilde{w}_i(h''_i)$ and there is no h'''_i such that $h'_i \not\subseteq h'''_i \not\subseteq h''_i$, then $A^q \cap A(h''_i) \neq \emptyset$
 - If $h'_i \not\subseteq h''_i$ 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 $\bar{h} \supseteq (h'_i, a)$, $i \in g_y(\bar{h})$.⁴⁹

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.

We now prove Theorem 6. As discussed in the main text, the first part of this theorem follows from our Corollary 1 and Li (2017) and the construction of a one-step simple strategic collection for each agent that we now present. This construction also proves the second part of the theorem.

Let Γ be a personal clock auction. We present the argument for “In-transfer falls”; the case for “Out-transfer falls” is analogous. Consider any $h_i \in \mathcal{H}_i$ and simple-node set $\mathcal{H}_{i,h_i} = \{h' \in \mathcal{H}_i : h_i \not\subseteq h' \not\subseteq h' \implies h' \notin \mathcal{H}_i\}$, and define the strategic plan $S_{i,h_i}(h')$ for any $h' \in \mathcal{H}_{i,h_i}$ as follows:

- If $\theta_i + \tilde{w}_i(h_i) > \bar{w}_i$ and $A(h_i) \setminus A^q \neq \emptyset$:
 - **[Action at h_i]** Choose $S_{i,h_i}(h_i) = a \in A(h_i) \setminus A^q$; if it further holds that $|A(h_i) \setminus A^q| > 1$, then choose $S_{i,h_i}(h_i) = a \in A(h_i) \setminus A^q$ such that $i \in g_y(\bar{h})$ for all $\bar{h} \supseteq (h_i, a)$.
 - **[Actions at next-histories]** For $h' \in \mathcal{H}_{i,h_i} \setminus \{h_i\}$, if there exists $a \in A(h') \cap A^q$, then $S_{i,h_i}(h') = a$ for some $a \in A(h') \cap A^q$. Otherwise, $S_{i,h_i}(h') = a'$ for some $a' \in A(h')$ such that for all $\bar{h} \supseteq (h, a')$, $i \in g_y(\bar{h})$.
- Otherwise, choose actions such that $S_{i,h_i}(h') \in A^q$ for all $h' \in \mathcal{H}_{i,h_i}$.

To show that this is a one-step simple strategic collection, consider first some h_i such that $\theta_i + \tilde{w}_i(h_i) > \bar{w}_i$ and there exists some $a \in A(h_i) \setminus A^q$. If $|A(h_i) \setminus A^q| = 1$, then all other

⁴⁹The corrigendum issued by Li replaces this statement with one that says if there is more than one non-quitting action at h'_i , there is a *continuation strategy* (rather than an action) that guarantees that $i \in g_y(\bar{h})$. The corrigendum also notes, though, that this change does not expand the set of implementable choice rules, because for any newly admissible mechanism, there is always an equivalent mechanism satisfying the original definition in which the agent reports her type at h'_i and does not move again. Thus, our notion of equivalence allows us to work directly with this simpler definition of personal clock auctions.

actions at h_i are quitting actions, and i 's best case (and worst case) payoff from following any such action is \bar{w}_i . We must show that the worst case from the perspective of node h_i from following the specified strategic plan gives a weakly greater payoff than \bar{w}_i . For any next-history $h'_i \in \mathcal{H}_{i,h_i}$ at which there is a quitting action (i.e., $A(h'_i) \cap A^q \neq \emptyset$), the worst case from the perspective of h_i of following the strategic plan is \bar{w}_i . If there is no quitting action at h'_i (i.e., $A(h'_i) \cap A^q = \emptyset$), then, by construction of a personal clock auction, we have that (i) $\tilde{w}_i(h_i) = \tilde{w}_i(h'_i)$, and (ii) there exists an $a' \in A(h'_i)$ such that, for all $\bar{h} \supseteq (h'_i, a')$, we have $i \in g_y(\bar{h})$. Further, for any $h''_i \not\supseteq h'_i$, $\tilde{w}_i(h''_i) = \tilde{w}_i(h'_i) = \tilde{w}_i(h_i)$, and so, for any $\bar{h} \supseteq (h'_i, a')$, $g_{w,i}(\bar{h}) = \tilde{w}_i(h_i)$. Therefore, the worst case from following the strategic plan from the perspective of h_i conditional on reaching any such h'_i is $\theta_i + \tilde{w}_i(h_i)$. In either case, the worst case from the strategic plan from the perspective of h_i is weakly better than taking any other action at h_i .

If $|A(h_i) \setminus A^q| > 1$, then the strategic plan instructs i to follow the action $a \in A(h_i)$ such that $i \in g_y(\bar{h})$ for all $\bar{h} \supseteq (h_i, a)$; further, by construction of a personal clock auction, at any $\bar{h} \supseteq (h_i, a)$, we have $g_{w,i}(\bar{h}) = \tilde{w}_i(h_i)$. Since $\theta_i + \tilde{w}_i(h_i) > \bar{w}_i$, this is strictly preferred to the payoff from taking any quitting action at h_i , and since the going transfer cannot rise, it is also weakly preferable to taking any other non-quitting action at h_i .

The last two cases two consider are when either (i) $A(h_i) \setminus A^q = \emptyset$ or (ii) $\theta_i + \tilde{w}_i(h_i) \leq \bar{w}_i$ holds. If (i) holds, then the only actions available at h_i are all quitting actions. Thus, the best- and worst-cases from any action are all \bar{w}_i , and it is trivial to see that one-step dominance holds. If (ii) holds, then, the worst case from quitting at h_i is a payoff of \bar{w}_i . Since the going transfer can only fall, the best case from playing a non-quitting action at h_i is at most $\theta_i + \tilde{w}_i(h_i) \leq \bar{w}_i$, and so again one-step dominance holds. ■

A.6 Proof of Theorem 7

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 $\bar{x}_{h^*} = \text{Top}(>_i, P_i(h^*))$ and $\bar{y}_{h^*} = \text{Top}(>_i, C_i(h^*))$. Let $\mathcal{H}_{i,h^*} = \{h \in \mathcal{H}_i | h^* \not\supseteq h' \not\supseteq h \implies h' \notin \mathcal{H}_i\}$ be the set of one-step 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 $S_{i,h^*}(h^*) = a_{\bar{x}_{h^*}}$.
 - Otherwise, we have $C_i(h^*) \subseteq C_i(h)$ (by monotonicity). In this case, $S_{i,h^*}(h^*) = a_{\bar{y}_{h^*}}$.

It is straightforward to verify that this strategic plan is one-step dominant at any h^* , and thus the corresponding strategic collection $(S_{i,h^*})_{h^* \in \mathcal{H}_i}$ is also one-step dominant.

Now, we prove the first statement. Let $(\Gamma, S_{\mathcal{N}, \mathcal{H}})$ be a millipede mechanism with a profile of one-step dominant strategic collections $S_{\mathcal{N}, \mathcal{H}}$. Begin by constructing an equivalent millipede mechanism that satisfies Lemma A.3. Note that the transformations used in the proof to construct the equivalent millipede mechanism are one-step dominance preserving—i.e., if $(\Gamma, S_{\mathcal{N}, \mathcal{H}})$ was an OSS millipede mechanism before the transformation, then the transformed game $(\Gamma', S'_{\mathcal{N}, \mathcal{H}})$ is another OSS millipede mechanism that satisfies Lemma A.3. We next show that any OSS millipede mechanism that satisfies Lemma A.3 is monotonic, which will complete the proof.

Let $(\Gamma, S_{\mathcal{N}, \mathcal{H}})$ be such a mechanism that is not monotonic, which means there exists an agent i , a history h^* at which i moves, another history $h \in \mathcal{H}_{i,h^*}$ such that $h \supseteq (h^*, a^*)$, 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 an earliest history at which monotonicity is violated in this way. Note that this implies that that $x \notin C_i(h')$ for any $h' \subseteq h^*$ such that $i_{h'} = i$, which in particular means that history h^* is on-path for any type \succ_i of agent i that ranks payoff x first.⁵⁰

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 $\succ_i: x, y, \dots$, and let \succ'_i be a type of agent such that $\succ'_i: x, z, \dots$. These two types exist by richness, given that we are in a no-transfer environment. Note that as argued in the previous paragraph, for $\succ''_i = \succ_i, \succ'_i$, we have $S_{i,h'}(\succ''_i)(h') = a^*$ for all $h' \subseteq h^*$.

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

Case (1): $y \notin P_i(h)$. From above, $S_{i,h^*}(\succ_i)(h^*) = a^*$. However, for any such strategic plan, the worst case outcome from the perspective of h^* is some $w \neq x, y$.⁵¹ Since she can clinch y at h^* , and $y \succ_i w$, $S_{i,h^*}(\succ_i)$ is not one-step dominant.

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

⁵⁰If $x \in C_i(h')$ for some h' , then, by monotonicity, at any next history $h'' \supseteq h'$ following a pass where i moves, either $x \in C_i(h'')$ or $P_i(h') \setminus C_i(h') \subseteq C_i(h'')$. If the latter holds, then at h'' , i has been offered to clinch everything that is possible for her, and so h is not on-path for any type of agent i , and we can construct an equivalent game in which monotonicity is not violated at h^* . Therefore, $x \in C_i(h'')$. Repeating this argument for every history between h' and h^* at which i moves delivers that $x \in C_i(h^*)$, which is a contradiction. Thus, x is possible at h^* , and therefore also possible at all $h' \subsetneq h^*$ at which i moves, but is not clinchable at any such h' , which means that any one-step dominant strategic plan must have $S_{i,h'}(\succ_i)(h') = a^*$ for any type \succ_i that ranks x first (and such a type exists by richness, given that we are in a no-transfer environment).

⁵¹If $x \notin P_i(h)$, 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 $y \notin P_i(h)$, $w \neq y$.

Subcase (2).(i): $z \in P_i((h, a^*))$. In this case, type \succ_i has no one-step dominant strategic plan at h^* . Again, in any such plan, we have $S_{i,h^*}(\succ_i)(h) = a^*$. 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 one-step dominant.

Subcase (2).(ii): $z \notin P_i((h, a^*))$. If $x \in P_i(h)$, then type \succ'_i has no one-step dominant strategic plan at h . To see this, note that at h , for any strategic plan, the worst case from passing at h is strictly worse than z ,⁵² while $z \in C_i(h)$, and so $S_{i,h}(h) \neq a^*$. However, $S_{i,h}(h)$ must equal a^* because x is possible, but not clinchable, at h ; a contradiction.

If $x \notin P_i(h)$, then type \succ_i has no one-step dominant strategic plan at h^* . 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 one-step dominant. ■

A.7 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 A.5. *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 A.5. 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

⁵²Note that here, we consider h , not h^* . The fact that x is not clinchable at h implies that x is not guaranteeable at h (recall the game satisfies Lemma A.3). This means that there must be some $h' \not\supseteq (h, a^*)$ such that i moves at h' or h' is terminal and such that $x \notin C_i(h')$. Combined with $z \notin P_i((h, a^*))$, this implies that the worst case from passing is strictly worse than z .

must be some type, \succ_i , such that action a is strongly obviously dominant at h_0^i . Without loss of generality, let $Top(\succ_i, P_i(h_0^i)) = x$. 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 be no better than z ; in particular, this implies that for all $a' \neq a$ and all $w \in P_i((h_0^i, a'))$, $w \not\geq y$. 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 $y \succ'_i w \succ'_i x$.⁵³ 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 $w \in P_i((h_0^i, a'))$, $w \not\geq x$. 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 once 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 $h_1 \not\subseteq h_2$ 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

⁵³If $x \geq w$ 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. Such types exist by richness and the fact that x, y, z are all mutually undominated.

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 $(h_1, a) \subset h_2$.

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 $Top(>_i, P_i(h_1, a_1^*)) = x$.

Next, we claim that for all $a' \neq a_1^*$ and all $w \in P_i((h_1, a'))$, we have $w \succeq y$. To see this, assume that there was some such a' and w such that $w \not\succeq y$. By the previous lemma, $w \not\succeq x$ for all $w \in P_i(h_1)$.⁵⁴ If $y \succeq w$, then $y \succ w$ (since $w \not\succeq y$). By pruning, some type $>'_i$ 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 $y \succ w$, we have $y >'_i w$, and so a' does not strongly obviously dominate a , which is a contradiction. If $y \not\succeq w$, then neither y nor w dominate each other, and type $x >_i w >_i y$ has no strongly obviously dominant action at h_1 . Therefore, $w \succeq y$ 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 $w \succ y$, then type $x >_i w >_i y$ again has no strongly obviously dominant action.⁵⁵ 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.⁵⁶ 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 $y' \succ y$. If $y \succ y'$, then type $x >_i y >_i y'$ has no strongly obviously dominant action at h_1 . Finally, to see that $x' = x$, again note that we cannot have $x' \succ x$; if $x \succ x'$, then type $x >_i y >_i x'$ has no strongly obviously dominant action at h_1 . Thus, $\hat{P}_i(h_2) = \{x, y\}$.

Finally, note that any type that prefers $x >_i y$ must select action a_1^* at h_1 , and thus, h_2 is on-path; further, any type that prefers $y >_i x$ must select some $a' \neq a_1^*$ at h_1 ; in other words, for all types that reach h_2 , $Top(>_i, P_i(h_2)) = x$. 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

⁵⁴If $w \succeq x$ for some $w \in P_i((h_1, a'))$, then, by the lemma, $\hat{P}_i((h_1, a')) = \{\hat{w}\}$, and $\hat{w} \succeq w \succeq x$, 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 $>_i$.

⁵⁵The 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^*).

⁵⁶If not, then all payoffs in $P_i(h_2)$ can be ordered by the dominance relation \succeq , and, if a strongly obviously dominant action exists, all types will take the same action, and the remaining actions can be pruned.

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 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 , $x \succeq_i x'$ 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 \succ_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 $Top(\succ_i, P_i(h_2)) = x$, a_2 is not strongly obviously dominant, which is a contradiction. ■

A.8 Proof of Theorem 8

That sequential choice mechanisms are SOSP is immediate from the definition, and so we focus on proving that every SOSP game is equivalent to a sequential choice mechanism. Note first that the pruning principle continues to apply to strong obvious dominance. Also, following the same reasoning as in the proof of Theorem 5, given any SOSP game, we can construct an equivalent SOSP 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 SOSP game in which there are no moves by Nature is equivalent to a sequential choice 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 5, 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 5 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 obvious 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 SOSP 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 choice mechanism. Assume not, and let h be an earliest history where the definition of a sequential choice mechanism is violated. Since Γ' is not a sequential choice 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 unclinchable x that is not dominated, i.e., $x \in \hat{P}_i(h)$.⁵⁷ 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 A.5 (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 $Top(>_i, P_i(h)) = x$ must select a , and $x \notin P_i((h, a'))$ for any $a' \neq a$.⁵⁸ Now, if $x \triangleright w'$, then type $x >_i w' >_i y$ 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 $x >_i w' >_i y$ 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 $x >_i w' >_i y$ 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

⁵⁷If 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.

⁵⁸Since x is not dominated at h , such a type exists by richness.

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 $x \succ_i y \succ_i z$ 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 $w \in P_i((h, a'))$, $x \triangleright y \succeq w$.⁵⁹ 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): $|P_i(h)| = 2$.⁶⁰

Let $P_i(h) = \{x, y\}$. Given the definition of a sequential choice 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 type will not have a strongly obviously dominant action. ■

A.9 Proof of Theorem 9

A.9.1 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.⁶¹ 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.⁶² 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

⁵⁹Note 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.)

⁶⁰If $|P_i(h)| = 1$, the argument is trivial.

⁶¹Our 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.

⁶²While 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.

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 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 .⁶³ 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 9 for symmetrized mechanisms.

Lemma A.6. *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 A.1, 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.

⁶³While 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 $S_{\sigma,i}(>_i)$ 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 $S_i^*(>_i)(h_A) = S_i^*(>_i)(h_B)$ for any equivalent histories h_A and h_B in these two games.

A.9.2 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 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 9.

A.9.3 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 A.3 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 $h' \not\supseteq h$.

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).⁶⁴ 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

⁶⁴They 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.

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 a 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 $C_i^{\subsetneq}(h) = \{x : x \in C_i(h') \text{ for some } h' \subsetneq h\}$ 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 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' .⁶⁵

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 .⁶⁶

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 A.13 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

⁶⁵This 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.

⁶⁶That 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.9.6 below.

$\ell_{m_1}^h$ 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 $y \in C_j^\varepsilon(h)$ for the other active non-lurker $j \in \bar{\mathcal{L}}(h) \setminus \{i\}$, then j is offered to clinch anything from what remains, $\bar{\mathcal{X}}^\mathcal{L}(h) \setminus \{y\}$. Note that such 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, $C_j^\varepsilon(h)$.

A.9.4 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_\emptyset to the terminal node \bar{h} and outputs a partial ordering of the agents, denoted \succ . 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, \succ_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 \succ 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 \succ 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 $C_j^c(h)$) 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_\emptyset 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 $\tilde{\succ}^k$ on the set of agents who are processed in step k . At the end of the final step K , we concatenate the K components to produce \succ , 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 .⁶⁷ 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 $y \in C_{j'}^c(h^1)$, 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

⁶⁷That 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 A.15, $i^1 \notin \mathcal{L}(h^1)$.

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 $\tilde{\succ}^k$.⁶⁸

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 \succ in the natural way: for any two agents i, j who were processed in the same step k , $i \succ j$ if and only if $i \tilde{\succ}^k j$. For any two agents i, j processed in different steps $k < k'$, respectively, we order $i \succ j$.

Remark. The output of the ordering algorithm is a partial order, \succ , on \mathcal{N} , the set of agents. If there are two agents j and j' such that $j \not\succeq j'$ and $j' \not\succeq j$, then we say j and j' **tie** under \succ . 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.9.5 Completing the proof

We complete the proof Theorem 9 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 \succ_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 \succ_σ if, for all j, j' : $j \succ_\sigma j'$ implies $f^{-1}(j) < f^{-1}(j')$. In other words, given some partial ordering \succ_σ , total order f is consistent if there is some possible way to break the ties in \succ_σ that delivers f .

Lemma A.7. *For any total order f consistent with \succ_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 $h_A^2 \not\preceq h_A^1$ is then the next time

⁶⁸At 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, $C_j^c(h^{k-1})$. 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 unclined objects after step $k - 1$. At the ‘‘root node’’ of this subgame, h_0^k , agent j is offered to clinch $C_j(h_0^k) \supseteq C_j^c(h^{k-1})$. All of the structure and arguments from the previous steps are then repeated.

along the game path that an agent who was not ordered in step 1 of the ordering algorithm clinches an object, etc.

Lemma A.8. *Let σ_A, σ_B be two role assignment functions, Γ_A and Γ_B their associated games, and $(\succ_A^1, \dots, \succ_A^{K_A})$ and $(\succ_B^1, \dots, \succ_B^{K_B})$ the respective partial orderings produced by each step of the ordering algorithm. For all k , if $\succ_A^k = \succ_B^k$, 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 A.8 implies the following corollary.

Corollary A.1. *If $\succ_A = \succ_B$, then $\sigma_A = \sigma_B$.*

Lemma A.9. *There are no three role assignment functions σ_A, σ_B and σ_C such that the resulting partial orders \succ_A, \succ_B and \succ_C take the form.⁶⁹*

$$\begin{aligned} i_1 \succ_A \cdots \succ_A i_n \succ_A \{i, j\} \cdots \\ i_1 \succ_B \cdots \succ_B i_n \succ_B i \succ_B j \cdots \\ i_1 \succ_C \cdots \succ_C i_n \succ_C j \succ_C i \cdots \end{aligned}$$

By Corollary A.1 of Lemma A.8, the mapping from role assignments σ to partial orders \succ_σ generated by the ordering algorithm is an injection. Lemma A.9 shows that we can break all the ties—recursively, coding step by coding step—creating from each \succ_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 A.7 because, by definition, each complete order f_σ generated by the tie-breaking in partial order \succ_σ is consistent with \succ_σ . ■

A.9.6 Preliminary Results for the Proofs of the Key Lemmas

Before proving the core Lemmas A.7, A.8, and A.9, we state several preliminary lemmas we will use. Lemmas A.10-A.14 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 A.3. Lemmas A.15 and A.16 are new.

⁶⁹What 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 \succ_A, \succ_B , and \succ_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.

Throughout, we fix a Pareto efficient millipede game Γ that satisfies properties (i)-(iii) of Lemma A.3.

Lemma A.10. (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 A.11. (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 A.12. (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,*

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 A.13. (BG Lemma E.19) *For all h , $|\bar{\mathcal{L}}(h)| \leq 2$.*

Lemma A.14. (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 $C_{i_h}(h) \cap C_t^c(h) = \emptyset$.*
- (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 $C_j^c(h') = \mathcal{X}(h)$.*

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

We also prove the following additional lemmas.

Lemma A.15. *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 A.15 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 $x_m^h \in C_j^c(h'')$, by BG Lemma A.11, 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 A.12, 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 .⁷⁰ 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 $h \subseteq h'' \not\subseteq h'$ and clinching 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 A.16. *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) $y \in C_i(h) \cap C_j^c(h)$, and define $\bar{x} = Top(\succ_j, \bar{\mathcal{X}}^{\mathcal{L}}(h))$. Then, $\bar{x} \succ_j y$.*

Proof. Let $h' \not\subseteq h$ be the largest subhistory such that $y \in C_j(h')$, and note that for this history, $P_j(h') = P_j(h)$.⁷¹ By construction, j passed at h' when she was offered to clinch y . If $P_j(h) \cap \bar{\mathcal{X}}^{\mathcal{L}}(h) = \emptyset$, then, by BG Lemma A.14, j is the terminator (i.e., $j = t$), and so by that same lemma, $C_i(h) \cap C_j^c(h) = \emptyset$, which is a contradiction (since $y \in C_i(h) \cap C_j^c(h)$). Therefore, $P_j(h) = \bar{\mathcal{X}}^{\mathcal{L}}(h)$,⁷² and so $P_j(h') = \bar{\mathcal{X}}^{\mathcal{L}}(h)$ as well. Since j passes at h' and $y \in C_j(h')$, $Top(\succ_j, P_j(h')) \succ_j y$. Since $P_j(h') = \bar{\mathcal{X}}^{\mathcal{L}}(h)$, $Top(\succ_j, P_j(h')) = Top(\succ_j, \bar{\mathcal{X}}^{\mathcal{L}}(h)) = \bar{x} \succ_j y$, as required. ■

A.9.7 Proofs of the Key Lemmas A.7, A.8, and A.9

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

⁷⁰Note 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.

⁷¹If $P_j(h) \not\subseteq P_j(h')$ (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 $h'' \not\subseteq h'$ where j was re-offered the opportunity to clinch y , by definition of a millipede game (or, more primitively, by OSP).

⁷²For 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).

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 .

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 2. 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 A.7

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 A.12, for each lurker ℓ_m , we have $x_m = \text{Top}(>_{\ell_m}, \mathcal{X} \setminus \{x_1, \dots, x_{m-1}\})$. 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 A.10, $\bar{\mathcal{X}}^{\mathcal{L}}(h^*) \subseteq P_i(h^*) \cup C_i^c(h^*)$, which implies that so $y = \text{Top}(>_{i^*}, \bar{\mathcal{X}}^{\mathcal{L}}(h^*))$. 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.⁷³ 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.⁷⁴ This process continues until someone eventually takes an unlurked object, all 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 $y \notin C_j^c(h^*)$, then the argument is exactly the same as in Case (1) (note that j is not ordered in step 1 in this case). If $y \in C_j^c(h^*)$, then the step 1 partial order is $\ell_1 \succ^1 \dots \succ^1 \ell_\lambda \succ^1 \{i^*, j\}$. 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 $h' \not\subseteq h^*$. Let $z = \text{Top}(\succ_j, \bar{\mathcal{X}}^\mathcal{L}(h^*))$, and note that by Lemma A.16, $z \succ_j y$. 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 $C_j^c(h^*)$ (note that $C_j^c(h^*) \subseteq \bar{\mathcal{X}}^\mathcal{L}(h)$). 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 $C_j^c(h^*)$ 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. ■

⁷³For 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.

⁷⁴When 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 $\Gamma_\sigma : \mathcal{X} \setminus \{x_1, \dots, x_{m'-1}\} \subseteq \mathcal{X} \setminus \{x_1, \dots, x_m\}$. 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.

Proof of Lemma A.8

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, $h_A^1 \not\subseteq h_B^1$. 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 $h_A^1 \not\subseteq h_B^1$, 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 \succ_A , let $g_A(i) = |j : j \succ_A i| + 1$.⁷⁵ Define g_B similarly. Since we assume $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$, 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 A.11). Since i_A clinches an unlurked object $x_A \in \bar{\mathcal{X}}^{\mathcal{L}}(h_A^1)$ at h_A^1 , we have $x_A = \text{Top}(\succ_{i_A}, \bar{\mathcal{X}}^{\mathcal{L}}(h_A^1))$,⁷⁶ 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 $\mathcal{X}^{\mathcal{L}}(h_A^1) = \mathcal{X}^{\mathcal{L}}(h_B^1)$.⁷⁷ 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).⁷⁸ 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 $\tilde{\succ}_B^1$, $x_A \in C_j^{\subseteq}(h_B^1)$, 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 $h' \not\subseteq h_A^1$ 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 A.16, $\bar{x} \succ_{i_A} x_A$. 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$), 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

⁷⁵This 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 \succ_A , i.e., $g_A(i) = g_A(i')$ if $i \not\subseteq_A i'$ and $i' \not\subseteq_A i$

⁷⁶This 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 A.10.

⁷⁷To see this, note that if $\mathcal{X}^{\mathcal{L}}(h_B^1) \not\subseteq \mathcal{X}^{\mathcal{L}}(h_A^1)$, 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 $h' \not\subseteq h_A^1$ 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 2).

⁷⁸Note that in this case, x_B is unlurked, and so all lurkers are immediately assigned to their lurked objects.

$g_A(i_A) = g_A(i_B) = g_B(i_A) = g_B(i_B) = \lambda(h_A^1) + 1$, i.e., i_A and i_B tie in both $\tilde{\succ}_A^1$ and $\tilde{\succ}_B^1$. 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 $h' \not\subseteq h_A^1$ 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 A.11. Further, the role/agent who moves at h_B^1 satisfies the conditions of the terminator t defined in BG Lemma A.14; denote $\rho(h_B^1) = t$, and note that $C_t^{\mathcal{C}}(h_B^1) = \mathcal{X}(h_B^1) = \mathcal{X}$. 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.

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 $y \in C_{i_A}^{\mathcal{C}}(h_B^1)$, 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$, 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 $x_A \in C_s^{\mathcal{C}}(h_A^1)$. But then, in game Γ_B , agent $i_A (= \sigma_B(s))$ was offered to clinch x_A at some history $h' \not\subseteq h_B^1$. 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 $x_A \notin C_t^{\mathcal{C}}(h_A^1)$, by BG Lemma A.14. 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 $\tilde{\succ}_A^1$), and so the same is true for g_B .⁷⁹ 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), $x_A \in C_{i_B}^{\mathcal{C}}(h_B^1)$, where $i_B = \sigma_B(r)$ (since $h_B^1 \not\supseteq h_A^1$ 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$ 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

⁷⁹Note 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 .

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

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

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 unlurked object y_A .⁸⁰ 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 A.1. 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$, 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 $\tilde{\succ}_A^1$ and $\tilde{\succ}_B^1$. 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, $x_{a_1}, x_{b_1} \in C_s^c(h^1)$.

If $\mathcal{L}(h^1) = \emptyset$, then $\mathcal{X} = P_{r_0}(h^1) \cup C_{r_0}^c(h^1)$, which implies that $x_{a_1} = \text{Top}(>_{i_A}, \mathcal{X})$ and $x_{b_1} = \text{Top}(>_{i_B}, \mathcal{X})$. But, this means that in Γ_A , i_B will clinch x_{b_1} at some $h' \not\subseteq h^1$ (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 A.10, $x_{a_1} = \text{Top}(>_{i_A}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1))$ and $x_{b_1} = \text{Top}(>_{i_B}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1))$. Since $x_{b_1} \in C_s^c(h^1)$, there is some history $h' \not\subseteq h^1$ such that in game Γ_A , $x_{b_1} \in C_{i_B}(h')$, and i_B passes at h' . By Lemma A.16, $\text{Top}(>_{i_B}, \bar{\mathcal{X}}^{\mathcal{L}}(h^1)) >_{i_B} x_{b_1}$, 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 A.14

⁸⁰Recall 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 $a_1 < a' < a_2$.

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 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 $\text{Top}(\succ_{i'}, \mathcal{X}) \neq y_A$). 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$. 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, $x_{a_M} = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_{M-1}}\})$. Similar to the previous paragraph, $\sigma_B^{-1}(i') \neq r_{n''}$ for any $n'' > a_M - 1$. Indeed, if this were true, then $x_{n''} = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{n''-1}\})$. If $n'' < a_M$, $x_{n''} \succ_{i'} x_{a_M}$, which contradicts $x_{a_M} = \text{Top}(\succ_{i'}, \mathcal{X} \setminus \{x_1, \dots, x_{a_{M-1}}\})$. If $n'' > a_M$, then x_{a_M} is not possible for i' (BG Lemma A.12), 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 $r' = \mathcal{A}_R(h^1) \setminus (\mathcal{L}_R(h^1) \cup \{r_0\})$.⁸¹ 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

⁸¹The case where $\sigma_B^{-1}(i') = r_0$ can be dispensed with similarly as in the previous paragraph.

$m = 1, \dots, M$. 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 $x_{a_1} = \text{Top}(\succ_{i'}, \mathcal{X})$. 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' . 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 $\tilde{\succ}_A^1 = \tilde{\succ}_B^1$, 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 .⁸² 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 $C_j^\Xi(h^1)$, if such an agent exists), and so we simply repeat the above arguments for each round $k = 1, \dots, K$. ■

Proof of Lemma A.9

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

As with Lemma A.8, we show this first for the case that under \succ_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.⁸³ In particular, this means that in \succ_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 A.15,

⁸²This 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 $h' \not\subseteq h^1$, 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.

⁸³That is, $h_\alpha^* = (h_\emptyset, a^*, \dots, a^*)$ for $\alpha = A, B, C$, though the number of passes (a^* 's) may vary.

and it is possible that $i_{h_A^*} = i$). The structure of \succ_A implies that $y \in C_j^c(h_A^*)$.

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

In this case, by definition of \succ_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 $y \in C_{s'}^c(h_A^*)$. Further, both i and j are getting their first choice objects, i.e., $Top(\succ_i, \mathcal{X}) = y$ and $Top(\succ_j, \mathcal{X}) = z$.

Now, consider σ_B , where the ordering algorithm produces $i \succ_B j$. By Remark 2, y cannot be the first lurked object in the game, and thus, for i to be ordered first without ties according to \succ_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 $j \succ_C i \dots$. 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 \succ_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 2, 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 $y \in C_{r'}^c(h')$ for the role $r' \in \{s, s'\}$ such that $\sigma_C^{-1}(j) \neq r'$. Therefore, i must clinch y at $h_C^* \not\supseteq h_A^*$. 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 $y \in C_k^c(h_C^*)$, and so, by construction of \succ_C , we have $j \succ_C \{i, k\} \succ_C \dots$, 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, $y \in C_k^c(h_A^*)$, which implies that i cannot be the next agent added to \succ_C uniquely—a contradiction.⁸⁴ The last case is $\sigma_C^{-1}(j) = s$. Note that $z \notin C_s^c(h_B^*)$,⁸⁵ and so $h_C^* \supset h_B^*$. This further implies that $y \in C_{s'}^c(h_C^*)$ 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 84, i cannot be the next agent added to the order uniquely, which is again a contradiction.

This completes the argument for Case 1.

⁸⁴For agent i to be ordered next in \succ_C without ties, she must be the first agent ordered in step 2 of the ordering algorithm. By Remark 2 again, y cannot be the next lurked object in the game, which means that i must clinch y at some $h' \not\supseteq h_C^*$ such that $\mathcal{L}(h') = \emptyset$. But, $y \in C_k^c(h')$ and so, by construction of \succ_C , she will tie with agent k .

⁸⁵If $z \in C_s^c(h_B^*)$, 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 84, j could not be the next agent added to \succ_B uniquely, a contradiction.

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).

Claim A.2. 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.⁸⁶ This implies that there is no passing action at h_A^* (by Lemma A.11) and that s is the terminator role t defined in BG Lemma A.14. Therefore, by BG Lemma A.14, $y \in C_s^{\subseteq}(h_A^*)$. That $y \in C_{s'}^{\subseteq}(h_A^*)$ for the other non-lurker role s' follows immediately from the construction of \triangleright_A . ■

Claim A.3. 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 $\mathcal{X}^{\mathcal{L}}(h_B^*) = \{x_1, \dots, x_{n'}\} \not\subseteq \mathcal{X}^{\mathcal{L}}(h_A^*) = \mathcal{X}^{\mathcal{L}}(h_B^*) \cup \{x_{n'+1}, \dots, x_n\}$. (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 A.11, 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 2. An analogous argument applies for h_C^* . ■

Claim A.4. $\mathcal{X}^{\mathcal{L}}(h_A^*) = \mathcal{X}^{\mathcal{L}}(h_B^*) = \mathcal{X}^{\mathcal{L}}(h_C^*)$.

Proof of claim. Given Claim A.3, 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 $h_B^* \not\supseteq h_A^*$. 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

⁸⁶By definition, $\sigma_A(s) \neq j$, and so this exhausts all possibilities.

$k \succ_B i$, which is a contradiction. However, since $h_B^* \not\supseteq h_A^*$, y has previously been offered to both active non-lurkers at h_B^* (from the construction of \succ_A). Thus, by Remark 2, 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} 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 $i_1 \succ_C \dots \succ_C i_n \succ_C j \succ_C i \dots$, 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 2, 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 84. Now, this implies that i was previously offered y at some $h'' \not\supseteq h_C^*$, 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 $\bar{x} = z$.⁸⁷ But, $z \notin \mathcal{X}^{\mathcal{L}}(h_A^*)$, which contradicts that she clinches y in Γ_A .⁸⁸ ■

Claim A.4 thus implies that $\mathcal{L}_R(h_A^*) = \mathcal{L}_R(h_B^*) = \mathcal{L}_R(h_C^*)$. Further, we know from BG Lemma A.13 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 $h_B^* \not\supseteq h_A^*$). A similar remark applies to h_C^* .

Claim A.5. $h_B^*, h_C^* \not\supseteq h_A^*$.

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 \succ_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

⁸⁷Each 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 $\bar{x} = z$.

⁸⁸There 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 .

j clinches z . Since $h_C^* \supseteq h_A^*$, both $s, s' \in \mathcal{A}_R(h_C^*)$, and $y \in C_s^c(h_C^*)$ and $y \in C_{s'}^c(h_C^*)$. Since i is the next agent ordered in \succ_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 84). Since i clinches $y \in \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*)$ in Γ_A , we have $y = \text{Top}(\succ_i, \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*))$. Since $\sigma_C(s') = i$, there is some $h' \not\subseteq h_A^*$ such that $y \in C_i(h')$ and i passes at h' in Γ_C . By Lemma A.16, $\text{Top}(\succ_i, \bar{\mathcal{X}}^{\mathcal{L}}(h_A^*)) = \bar{x} \succ_i y$, which is a contradiction. ■

Claim A.6. 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 A.5, $h_B^* \not\subseteq h_A^*$. 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 A.11. An equivalent argument holds for h_C^* . ■

In words, Claim A.6 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 A.9. 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 $y \in C_{s'}^c(h_A^*)$. Now, consider σ_B/\succ_B . By the discussion following Claim A.6, $i_{h_B^*} = i$. This implies that $\sigma_B^{-1}(i) = s'$.⁸⁹

Now consider σ_C/\succ_C . Again, by the discussion following Claim A.6, $i_{h_C^*} = j$; further, $\sigma_C^{-1}(j) = s$.⁹⁰ If $z \in C_s^c(h_B^*)$, then, by an argument equivalent to footnote 84, it must be that $\sigma_B(s) = j$. Further, $h_B^* \not\subseteq h_C^*$ (otherwise, j would clinch at $h_C^* \not\subseteq h_B^*$ in Γ_B , since she has the same role in both games). Therefore, $y \in C_{s'}^c(h_C^*)$ (in particular, $y \in C_{s'}(h_B^*)$). Again by the same argument as footnote 84, $\sigma_C(s') = i$. But, $\sigma_C(s') = i$ implies that i clinches at $h_B^* \not\subseteq h_C^*$ in Γ_C (since i has the same role as in Γ_B), which is a contradiction. Finally, if $z \notin C_s^c(h_B^*)$, we once again have $h_B^* \not\subseteq h_C^*$ and so $y \in C_{s'}^c(h_C^*)$, 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

⁸⁹In particular, we cannot have $\sigma_B^{-1}(i) = s$ again, because this would imply $h_B^* = h_A^*$, contradicting Claim A.5.

⁹⁰If $\sigma_C^{-1}(j) = s'$, then, j must pass at all $h' \not\subseteq h_A^*$ 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 $h_C^* \not\subseteq h_A^*$, which contradicts Claim A.5.

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 A.11. Further, this implies that role s is the terminator role defined in BG Lemma A.14.

In game Γ_B , the discussion following Claim A.6 again gives $i_{h_B^*} = i$. If $\sigma_B^{-1}(i) = s$, then, 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 A.6 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 A.14, 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 $h' \not\subseteq h_B^*$ be a history such that $z \in C_s(h')$. By an argument equivalent to footnote 84, 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 $h_C^* \not\supseteq h_B^*$, and so $y \in C_{s'}^c(h_C^*)$ (in particular, $y \in C_{s'}(h_B^*)$). By an argument equivalent to footnote 84, $\sigma_C(s') = i$. However, if $\sigma_C(s') = i$, then i clinches at $h_B^* \not\subseteq h_C^*$ 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 $\succ_A, \succ_B, \succ_C$ of the form given in the statement of Lemma A.9 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 A.9. ■