A Theory of Simplicity in Games and Mechanism Design

Marek Pycia and Peter Troyan*

February 2023

Abstract

We study extensive-form games and mechanisms allowing agents that plan for only a subset of future decisions they may be called to make (the *planning horizon*). Agents may update their so-called *strategic plan* as the game progresses and new decision points enter their planning horizon. We introduce a family of *simplicity* standards which require that the prescribed action leads to unambiguously better outcomes, no matter what happens outside the planning horizon. We employ these standards to explore the trade-off between simplicity and other objectives, to characterize simple mechanisms in a wide range of economic environments, and to delineate the simplicity of common mechanisms such as posted prices and ascending auctions, with the former being simpler than the latter.

*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." The accepted paper contained a characterization of Random Priority as the unique obviously strategy-proof, Pareto efficient, and symmetric mechanism. Due to length concerns, we were asked to shorten the paper after acceptance. As the characterization was essentially independent from the rest of the paper and took nearly half of its length, we proposed removing the characterization as a way to meet the length constraints, and the editor agreed to its removal. For their comments, we would like to thank Itai Ashlagi, Sarah Auster, Eduardo Azevedo, Roland Benabou, Dirk Bergemann, Tilman Börgers, Ernst Fehr, Dino Gerardi, Ben Golub, Yannai Gonczarowski, Ed Green, Samuel Haefner, Rustamdjan Hakimov, Shaowei Ke, Fuhito Kojima, Simon Lazarus, Jiangtao Li, Shengwu Li, Giorgio Martini, Stephen Morris, Anna Myjak-Pycia, Nick Netzer, Ryan Oprea, Ran Shorrer, Erling Skancke, Susan Spilecki, Utku Ünver, Roberto Weber, anonymous referees, the Eco 514 students at Princeton, and the audiences at the 2016 NBER Market Design workshop, NEGT'16, NC State, ITAM, NSF/CEME Decentralization, the Econometric Society Meetings, UBC, the Workshop on Game Theory at NUS, UVa, ASSA, MATCH-UP, EC'19 (the Best Paper prize), ESSET, Wash U, Maryland, Warsaw Economic Seminars, ISI Delhi, Notre Dame, UCSD, Columbia, Rochester, Brown, Glasgow, Singapore Management University, Matching in Practice, Essex, European Meeting on Game Theory, GMU, Richmond Fed, Israel Theory Seminar, USC, and Collegio Carlo Alberto. 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.

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No 866376.

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 an extensive-form game, with the final outcome determined by the decisions made by the agents during the game. To ensure that the final outcome satisfies desirable normative properties (e.g., efficiency or revenue maximization), the standard approach in mechanism design is to provide agents with incentives to play in a predictable and desirable way. For instance, the designer may use a Bayesian or dominant-strategy incentive-compatible direct mechanism where it is in each agent's best interest to simply report all of their private information truthfully. This approach succeeds so long as the participants understand that being truthful is in their interest (for instance, if the designer has the ability to successfully teach the agents how to play). However, there is accumulating empirical evidence of agents not reporting the truth in such mechanisms.[1] Bayesian or dominant-strategy mechanisms are thus sometimes not sufficiently simple for participants to play as expected. Using simpler mechanisms can reduce such strategic confusion. It may also 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 for games and mechanism design. We use these standards to assess the restrictions that simplicity imposes on the mechanism designer and to characterize simple mechanisms for a broad range of social choice environments with and without transfers.[3] Analogously to the revelation principle for Bayesian mechanism design, we construct classes of mechanisms that limit the space over which a designer interested in implementing a simple mechanism must search.[4] As applications, we provide simplicity-based microfoundations for popular mechanisms such as posted prices, priority mechanisms, and ascending auctions.

[1]See, 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).

[2]See 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.

[3]Examples 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).

[4]Direct mechanisms are not necessarily simple, and hence 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 have unlimited foresight and are able to plan a complete strategy for every possible future contingency. Rather, we allow for agents with limited foresight who, each time they are called to play, make plans for only a subset of possible future moves, their current *planning horizon*. 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 decision points not planned for. As the game progresses, agents may update their strategic plans, and choose an action that is different from what they planned in the past. This potential for updating is what differentiates strategic plans from the standard game-theoretic concept of a strategy.[6]

We model variations in simplicity and required foresight ability by varying agents' planning horizons. This gives rise to a family of simple dominance standards. The stronger the simplicity standard—i.e., the fewer information sets in the current planning horizon—the more robust the corresponding mechanisms is to agents who can plan for only limited future horizons.[7] We show that the longer the planning horizon of the agents, the more social choice rules a designer can implement in a simply-dominant way. Furthermore, any implementable social rule can be implemented via a perfect-information extensive-form game.

We focus on special cases of simple dominance in which agents are able to plan some exogenously given number $k \in \{0, 1, ..., \infty\}$ of future moves and analyze three such special cases of simple dominance in detail.

- $k = \infty$: each agent's planning horizon consists of all information sets at which this agent moves (and contains no other information sets); 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. In this case, simple dominance becomes equivalent to Li's (2017b) obvious dominance, and so we refer to the resulting simply dominant strategic plans as *obviously dominant*, and the corresponding mechanisms as *obviously strategy-proof (OSP).*

- $k = 1$: each agent's planning horizon consists of their current information set and only

[5] Savage (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.

[6] We 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. The future actions in the partial strategic plan merely ensure the optimality of the initial action.

[7] We 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.

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)*.

- $k = 0$: each agent's planning horizon consists only of their current information set; in other words, agents cannot plan for any moves in the future. We refer to the resulting resulting simply dominant strategic plans as *strongly obviously dominant*, and the corresponding mechanisms as *strongly obviously strategy-proof (SOSP)*.

The above standards are nested: strongly obviously dominant strategic plans are one-step dominant, which in turn are obviously dominant. Obvious dominance is the most permissive of these standards. 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). For instance, 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. Winning strategies in chess require looking many steps into the future, and thus are not one-step dominant dominant nor strongly obviously dominant. 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'. Clinching corresponds to receiving a payoff for sure, and leaving the game. The agent may also be given the opportunity to 'pass'. If the agent passes, she remains in the game, with the potential of being offered better clinching options in the future. Agents are sequentially presented with such clinch-or-pass choices until all agents' payoffs are determined. The millipede class includes as special cases some familiar, and intuitively simple, games, such as serial dictatorships. However, the millipede class also admits other games that are rarely observed in market-design practice, and whose 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).

We next study one-step dominance in environments both with and without transfers. We 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 OSP mechanisms are equivalent to personal clock auctions. In particular, 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 OSP millipede games discussed above (and also eliminates games such as chess). Indeed, we characterize OSS millipede games as those that satisfy a property we call monotonicity. Monotonicity ensures that, at the first time an agent may be called to play again in the continuation game that begins following any of her current actions, the agent's best clinching option will be weakly better than anything she could have clinched in the past. Monotonic games seem particularly simple, 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 SOSP 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. Strongly obviously dominant strategic plans are incentive-compatible even for agents concerned about trembles, or who have time-inconsistent preferences. Building on this insight, we show that all SOSP mechanisms can be implemented as *sequential choice mechanisms* 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, 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 strongly obviously dominant choice. The offered menu may include prices, in which case we call the mechanism a *sequential posted price mechanism.* The strong obvious dominance of these mechanisms provides an explanation of the popularity of posted prices, a ubiquitous sales procedure.[9]

Our construction of 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 in two ways. First, we introduce gradated standards of simplicity, which allow us to assess the trade-off between simplicity and implementation flexibility. Second, we

[8]Monotonicity is a generalization of a similar feature of ascending auctions: in an ascending auction, if an agent passes (continues in the auction), at any next move, she is able to drop out (clinch the zero payoff), except if she wins. Since the zero payoff is the only clinchable payoff and this payoff is always clinchable (except if the agent wins), ascending auctions are monotonic. Such monotonicity is also satisfied by Li's personal clock auctions.

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

provide simplicity-based microfoundations for popular mechanisms such as posted prices and priority mechanisms. 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) mechanisms, and provides a characterization of the priority structures under which TTC are OSP-implementable.[10] Following our work, Arribillaga et al. (2020) and Arribillaga et al. (2019) characterize the voting rules that are obviously strategy-proof on domains of single-peaked preferences. Bade and Gonczarowski (2017) study obviously strategy-proof and efficient social choice rules in several environments. Mackenzie (2020) 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.[11]

Our work also contributes to the understanding of limited foresight and limits on backward induction. Other work in this area, with different approaches from ours, includes Jehiel (1995; 2001) on limited foresight equilibrium in which players' forecasts are correct, Gabaix et al. (2006) on directed cognition, and Ke's (2019) axiomatization of bounded-horizon backward induction. A major difficulty for models of imperfect foresight is how an agent takes into account a 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.[12] As a by-product, our mechanisms work well when agents are inter-temporally inconsistent, for instance because they face Knightian uncertainty or optimize against multiple priors (as, e.g., in Knight (1921) and Bewley (1987)) or have time-inconsistent preferences (as, e.g., in

[10] Li showed that the classic 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 priority structures under which Hierarchical Exchange of Pápai (2000) and Trading Cycles (group strategy-proof and efficient mechanisms) of Pycia and Ünver (2017) are OSP-implementable; cf. also Mandal and Roy (2021).

[11] Also of note is Glazer and Rubinstein (1996), who argued that extensive-form games may simplify the solution of normal-form games, and 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).

[12] The 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).

Strotz (1956) and Laibson (1997)).

Finally, this paper 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 Ehlers (2002) and Pycia and Unver (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.[13] Liu and Pycia (2011), Pycia (2011), Morrill (2015), Hakimov and Kesten (2014), Ehlers and Morrill (2017), and Troyan et al. (2020) characterize mechanisms that satisfy incentive, efficiency, and fairness objectives.

2 Model

2.1 Preferences

Let $\mathcal{N} = \{i_1, \ldots, i_N\}$ be a set of agents, and $\mathcal{X}$ a finite set of outcomes.[14] An outcome might involve a monetary transfer. Each agent has a preference ranking over outcomes, where, for $x, y \in \mathcal{X}$, we write $x \succsim_i y$ to denote that x is weakly preferred to y. We allow for indifferences, and write $x \sim_i y$ if $x \succsim_i y$ and $y \succsim_i x$. For any $\succsim_i$, we let $\succ_i$ denote the corresponding strict preference relation, i.e., $x \succ_i y$ if $x \succsim_i y$ but not $y \succsim_i x$. We use $\mathcal{P}_i$ to denote the domain of agent i's preferences, and refer to $\succsim_i$ 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.[15]

[13] Pycia 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) and Root and Ahn (2020).

[14] Assuming $\mathcal{X}$ 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.

[15] It is natural to assume that an agent weakly prefers lottery μ over ν whenever for all outcomes $x \in \operatorname{supp}(\mu)$ and $y \in \operatorname{supp}(\nu)$ 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.

2.2 Extensive-Form Games

To determine the outcome the planner designs a game Γ for the agents to play. We consider imperfect-information, extensive-form games with perfect recall. These 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 are 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 chooses an **action** from a finite set $A(h)$. 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. When Nature moves, she selects an action from $A(h)$ according to some known probability distribution.

The set of histories at which agent i moves is denoted $\mathcal{H}_i = \{h \in \mathcal{H} : i_h = i\}$. We partition $\mathcal{H}_i$ into **information sets** and denote this partition by $\mathcal{I}_i$. 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 $h_1 \subseteq h_2$. If I_1 precedes I_2, we write $I_1 \leq I_2$ (and $I_1 < I_2$ if $I_1 \neq I_2$); we then also say that I_2 **follows** I_1 and that I_2 is a **continuation** of I_1. 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

What extensive-form games are simple to play? Intuitively, choosing from a menu of outcomes—e.g., a take-it-or-leave-it opportunity to buy an object at a posted price, or a choice from a set of objects in an extensive-form serial dictatorship (cf. Section 4.3)—entails simpler strategic considerations than those faced by a bidder in an ascending clock auction. Similarly, ascending clock auctions are simpler than complex games such as chess. We propose a

While our results do not rely on this assumption, it ensures that dominant actions always lead to weakly preferred lotteries over outcomes.

class of simplicity standards that allows us to differentiate the strategic simplicity of these and other games.

In defining our class of simplicity standards, we relax the standard assumption of economic analysis that players can analyze and plan their actions arbitrarily far into the future. Such unlimited 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 agents who do not reason carefully about what their opponents will do, while still requiring that they search deep into the game with regard to their future self. Li's agents know the structure of precedence among the information sets at which they move and the sets of outcomes that could possibly obtain conditional on any sequence of their own actions (though not conditional on their opponents' actions). For instance, if White has a winning strategy in chess—i.e., at the start of the game, White knows what to do at any possible future configuration of the board to ensure a victory—then this strategy is also obviously dominant. We relax Li's foresight assumptions, only maintaining that players know possible outcomes of actions and precedence relations for information sets in their planning horizon (i.e., those information sets for which the agent plans).

The key innovation in our framework is that an agent may update their plan as the game is played. In other words, we allow the agent's perception of the strategic situation, and hence, their planned actions to vary as the game progresses. To differentiate them from the standard game-theoretic notion of a "strategy" as a complete contingent plan of action, we refer to these objects as "strategic plans", introduced in the next subsection.

3.1 Strategic plans

Formally, each information set $I^* \in \mathcal{I}_i$ at which agent i moves has an associated set of continuation information sets $\mathcal{I}_{i,I^*} \subseteq \{I \in \mathcal{I}_i | I \geq I^*\}$ that are **simple from the perspective of** I^*; we call $\mathcal{I}_{i,I^*}$ agent i's **planning horizon** at I^*. We assume that $I^* \in \mathcal{I}_{i,I^*}$, but otherwise, the only restriction is that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}_i$.[16] A **(partial) strategic plan** $S_{i,I^*}(\succ_i)$ for agent i of type $\succ_i$ at information set I^* maps each simple information set $I \in \mathcal{I}_{i,I^*}$ to an

[16]The assumption that $\mathcal{I}_{i,I^*} \subseteq \mathcal{I}_i$ is what makes our simplicity standards dominance standards. Dropping this assumption and endowing players with beliefs of what other players do at simple information sets leads to an analogue of our theory for equilibria. 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.

action in $A(I)$.[17] Note that a strategic plan does not specify the play at *all* continuation information sets at which i may be called to move; rather, the strategic plan at I^* only specifies an action at the information sets in the planning horizon at I^*. Sets of strategic plans $(S_{i,I^*}(\succ_i))_{I^*\in\mathcal{I}_i}$ and $(S_{i,I^*}(\succ_i))_{I^*\in\mathcal{I}_i,\succ_i\in\mathcal{P}_i}$ 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, $S_{\mathcal{N},\mathcal{I}} = ((S_{i,I^*}(\succ_i))_{I^*\in\mathcal{I}_i,\succ_i\in\mathcal{P}_i})_{i\in\mathcal{N}}$. For any strategic collection $(S_{i,I^*}(\succ_i))_{I^*\in\mathcal{I}_i}$, we define the **induced strategy** $\hat{S}_i(\succ_i): \mathcal{I}_i \to \cup_{I\in\mathcal{I}_i} A(I)$ as the mapping from information sets to actions defined by $\hat{S}_i(\succ_i)(I) = S_{i,I}(\succ_i)(I)$ for each $I \in \mathcal{I}_i$; that is, $\hat{S}_i(\succ_i)$ 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 $\succ_{\mathcal{N}}$, we can determine the terminal history and associated outcome that is reached when the game is played according to the profile of strategic collections $S_{\mathcal{N},\mathcal{I}}(\succ_{\mathcal{N}})$ by following the profile of induced strategies $\hat{S}_{\mathcal{N}}(\succ_{\mathcal{N}})$. For each player i and type $\succ_i$, the induced strategy $\hat{S}_i(\succ_i)$ also allows us to define the set of **on-path information sets** for a strategic collection. These are the information sets $I \in \mathcal{I}_i$ such that there exists strategies for the other players and Nature such that I is on the path of play of $\hat{S}_i(\succ_i)$.

Induced strategies 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 $\succ_{\mathcal{N}}$, the distribution over outcomes from the induced strategies $\hat{S}_{\mathcal{N}}(\succ_{\mathcal{N}})$ in Γ is the same as that from the induced strategies $\hat{S}'_{\mathcal{N}}(\succ_{\mathcal{N}})$ in Γ'. 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 planning horizons for the agents. Every mechanism implements a mapping from preference profiles to outcomes, which we call the **social choice rule**. If two mechanisms are equivalent, they implement the same social choice rule.

3.2 Simple dominance

Strategic plan $S_{i,I^*}(\succ_i)$ is **simply dominant at information set** I^* for type $\succ_i$ of player i if the worst possible outcome for i in the continuation game assuming i follows $S_{i,I^*}(\succ_i)(I)$ 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 $a' \neq S_{i,I^*}(\succ_i)(I^*)$ at I^*. (We provide an example below to illustrate this definition.) We say that a strategic collection $(S_{i,I^*}(\succ_i))_{I^*\in\mathcal{I}_i,\succ_i\in\mathcal{P}_i}$ is **simply dominant** if, for each type $\succ_i\in\mathcal{P}_i$, the strategic plan $S_{i,I^*}(\succ_i)$ is simply dominant at I^* for each on-path information set I^*.[18] We say that a game is simply dominant if it admits

[17] We focus on pure strategies; the extension to mixed strategies is straightforward.

[18] When assessing $S_{i,I^*}(\succ_i)(I)$, we take the worst case over all game paths consistent with i following $S_{i,I^*}(\succ_i)(I)$ at all $I \in \mathcal{I}_{i,I^*}$, and compare to the best case over *all* game paths following any alternative

simply dominant strategies.

Note that the collection of planning horizons, $(\mathcal{I}_{i,I^*})_{I^* \in \mathcal{I}_i}$, is a parameter of the model. In the sequel, we focus on planning horizons that vary agents' length of foresight. This is not necessary, however, and there are other ways to conceptualize what information sets are in the planning horizons.[19] Given a fixed $k \in \{0, 1, 2, ..., \infty\}$, we say that agent i has **k-step foresight** if

$$\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I \text{ and } I^* < I_1 ... < I_k < I \Rightarrow \exists \ell \in \{1, ..., k\} \text{ s.t. } I_\ell \notin \mathcal{I}_i\}.$$

We refer to the resulting simply dominant strategic collections as **k-step dominant** and 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, $\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I\}$, and i can plan all of her future moves. In this case, the induced strategy of any resulting strategic collection is obviously dominant in the sense of Li (2017b), and further, any obviously dominant strategy S_i in the sense of Li (2017b) determines an obviously dominant strategic collection $(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ in which $S_{i,I^*}(I) = S_i(I)$ for any $I^* \leq I$. For these reasons, we refer to such strategic collections as **obviously dominant**. If a mechanism admits obviously dominant strategic collections, then we say it is **obviously strategy-proof (OSP)**.

- $k = 1$: That is, $\mathcal{I}_{i,I^*} = \{I \in \mathcal{I}_i | I^* \leq I \text{ and } I^* < I' < I \Rightarrow I' \notin \mathcal{I}_i\}$, and 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 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 $\mathcal{I}_{i,I^*} = \{I^*\}$, and 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

action $a' \neq S_{i,I^*}(\succ_i)(I^*)$. 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 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.

[19]For instance, the planning horizon could consist of the information sets 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.

$(S_{i,I^*})_{I^* \in \mathcal{I}_i}$ 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)**.

Example. (Simple Dominance in Ascending Auctions). We illustrate simple dominance by looking at an ascending auction for a single good. The payoff of an agent i is equal to the agent's value v_i minus their payment if the agent receives the good and it is equal to minus their payment otherwise. For the purposes of this example, an **ascending auction** is a finite game with the following properties. At each non-terminal information set an agent is called to play and can take one of two possible actions: Stay In or Drop Out. Only agents who have not dropped out yet—called active agents—are called to play. Each non-terminal information set I is associated with the current price $p(I) \geq 0$, which weakly increases along each path of play. Whenever there is only one active agent left, the game ends, though it can also end when there are several active agents. One of the agents active when the game ends is designated the winner. The winner receives the good and pays the price associated with the last history at which this agent moved; a winner who has not yet moved pays 0. All other agents receive no good and pay 0.

The ascending auction is OSP because the strategy of staying in as long as the current price is below the agent's value is obviously dominant, as shown by Li (2017b). The ascending auction is also OSS because the following strategic collection is one-step dominant. For any information set I^* at which i moves and $p(I^*) \leq v_i$, i's strategic plan is $S_{i,I^*}(I^*) = In$ and $S_{i,I^*}(I) = Out$ for all next information sets $I > I^*$; for any information set I^* at which i moves and $p(I^*) > v_i$, the strategic plan is $S_{i,I^*}(I) = Out$ for $I = I^*$ and for all next information sets $I > I^*$. This strategic collection is one-step dominant because if $p(I^*) \leq v_i$ then staying in at I^* and planning to drop out at the next information set gives a worst-case payoff of 0, which is no worse than dropping out now; if $p(I^*) > v_i$ then dropping out at I^* is weakly better than any scenario following staying in. The ascending auction might be SOSP, for instance when the starting prices are higher than the agents' values. In general, however, the ascending auction is not SOSP: Let i be the first mover and suppose the prices i might see along the path of the game start strictly below v_i, but prices further along the path are strictly above v_i. Then, at the first move of i, no move is strongly obviously dominant: the worst case from choosing In results in a strictly negative payoff, which is worse than choosing Out and getting 0. Thus, In is not strongly obviously dominant. Similarly, the payoff from choosing Out is 0, which is worse than choosing In and winning at the current low price, and so Out is not strongly obviously dominant, either.

Remark 1. *(Plan updating and consistency)* In the one-step dominant strategic collections in the example above, the action an agent plans at I^* for the next information set $I >$

I^* may differ from the action the agent chooses upon actually reaching I, i.e., we may have $S_{i,I^*}(I) \neq S_{i,I}(I)$. There is no need for such action updating in obviously dominant or strongly obviously dominant strategic collections. For each OSP collection there is an equivalent OSP collection that is *consistent* in the following sense: $S_{i,I^*}(I) = S_{i,I}(I)$ for all $I \in \mathcal{I}_{i,I^*}$ and all $I^* \in \mathcal{I}_i$. SOSP collections are always consistent.

We emphasize that agents with inconsistent strategic plans are not necessarily time-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).

3.3 Simplicity Gradations and Design Flexibility

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

Theorem 1. ***(Nesting of Simplicity Concepts).*** *If planning horizons $\mathcal{I}_{i,I^*}$ 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 $k, k' \in \{0, 1, 2, ..., \infty\}$ 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.[20]

Theorem 2. ***(Simplicity-Flexibility Tradeoff)*** *Let $k, k' \in \{0, 1, 2, ..., \infty\}$ and assume $k' > k$. There exist preference environments and social choice rules implementable in k'-step simple strategic collections, but not implementable in k-step simple strategic collections.*

The presence of the simplicity-flexibility 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 agent's value for an object comes from the same support with at least three distinct values. Suppose we want to allocate the 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, the comparison is more subtle. 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.

3.4 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 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 is unaffected by the player's confusion or partial description of the game.

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

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

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 to 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 $\succ_i$, 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^*.*

When the strategic collection is consistent, this result says that the induced global strategy $S_i(I) = S_{i,I}(I)$ is simply dominant in one game if and only if $S_i(I)$ is weakly dominant in all indistinguishable games. When expressed in this way, this result corresponds to Li's (2017b) microfoundation for obvious strategy-proofness.[21]

[21] While 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 while we only require that λ is an injection. This difference has no impact on the validity of the claim nor the proofs.

3.5 Design Sufficiency of Perfect-Information Games

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

Perfect-information games play a special role in designing simply dominant mechanisms because for any imperfect-information simply dominant mechanism, we can find an equivalent perfect-information one.[22] 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 $\left(\hat{S}_{i,h^*}(\succ_i)\right)_{h^* \in \mathcal{H}_i}$ in Γ', then the corresponding information set $I \ni h$ is on-path for the strategic collection $\left(\hat{S}_{i,I^*}(\succ_i)\right)_{I^* \in \mathcal{I}_i}$ 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 $a \neq S_{i,h^*}(h)$ over the entire information set h. 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'.

In light of Theorem 4, the restriction to perfect information games does not affect the class of social choice rules that can be implemented in simple strategies. We hence adapt this restriction in the study of mechanism design in the next two sections.

[22] An analogous property of obvious strategy-proofness was first asserted in a footnote in Ashlagi and Gonczarowski (2018). Following our work, Mackenzie (2020) extended this property of obvious strategy-proofness to extensive-form games without perfect recall.

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 allow some restrictions on the domains of agent preferences. We formalize this as follows.

We take as a primitive a **structural dominance relation** over outcomes, denoted $\trianglerighteq$, where $\trianglerighteq$ is a reflexive and transitive binary relation on $\mathcal{X}$. The notation $x \trianglerighteq y$ is read as "x weakly dominates y".[23] If $x \trianglerighteq y$ but not $y \trianglerighteq x$, then we write $x \triangleright y$, and say that "x strictly dominates y". For instance, in environments with transfers, outcome x dominates 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 $\trianglerighteq$ if $x \trianglerighteq y$ implies that $x \succsim_i y$ and $x \triangleright y$ implies that $x \succ_i y$.

We allow the possibility that different agents have different dominance relations, $\trianglerighteq_i$, and therefore different preference domains. We assume that all rankings in $\mathcal{P}_i$ are consistent with $\trianglerighteq_i$. If $x \trianglerighteq_i y$ and $y \trianglerighteq_i x$ then x and y are $\trianglerighteq_i$-**equivalent**. Any $\trianglerighteq_i$ determines an **equivalence partition** of $\mathcal{X}$. We refer to each element $[x]_i = \{y \in \mathcal{X} : x \trianglerighteq_i y \text{ and } y \trianglerighteq_i x\}$ of the equivalence partition as a **payoff**. 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 phrases such as "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]

The framework of rich preference domains is flexible and encompasses many standard

[23]For brevity, we write "weakly dominates" rather than "weakly structurally dominates" when the context makes clear that we refer to outcomes, comparable in terms of $\trianglerighteq$, and not to strategies, comparable in the game-theoretic sense of weak dominance.

[24]Our 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 $\trianglerighteq_i$, the smaller the resulting preference domain and less restrictive the simple dominance requirement. At one extreme, $\trianglerighteq_i$ is an identity relation for each $i \in \mathcal{N}$, agents' preference domains consist of all strict rankings, and simple mechanisms resemble dictatorships as in Gibbard (1973) and Satterthwaite (1975) and our Corollary 2. At the other extreme, $\trianglerighteq_i$ compares all outcomes, each agent is indifferent among all outcomes, and any strategy in any game is 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.

economic environments. Some examples of rich domains will help clarify the definitions and notation.

- **Voting:** Every agent has strict preferences over all alternatives in $\mathcal{X}$. This is captured by the trivial dominance relation $\trianglerighteq_i$ in which $x \trianglerighteq_i y$ implies $x = y$ for all i. Each agent's preference domain $\mathcal{P}_i$ partitions $\mathcal{X}$ into $|\mathcal{X}|$ individual payoffs. 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 $\trianglerighteq_i$ for agent i defined as follows: $x \trianglerighteq_i y$ if and only if agent i receives the same set of goods in outcomes x and y. Each payoff of agent i 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 these two examples in mind, we say that an environment is **without transfers** if the dominance relation $\trianglerighteq_i$ is symmetric for all i.[25] Non-symmetric dominance relations $\trianglerighteq_i$ allow 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 $\mathcal{W} \subsetneq \mathbb{R}$ a (finite) set of possible transfers. For a fixed $y \in \mathcal{Y}$, agent i prefers to pay less rather than more and is indifferent between any two outcomes that vary only in other agents' transfers. The structural dominance relation is $(y, w) \trianglerighteq_i (y', w')$ 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 $\mathcal{W} \subsetneq \mathbb{R}$ 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 $(O; w) \trianglerighteq_i (O'; w')$ if and only if $O_i \supseteq O_i'$ and $w_i \geq w_i'$.

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. Arribillaga, Massó, and Neme (2020) show that our millipede construction does not extend to single-peaked preference domains.

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

4.1 Obvious Dominance

By Theorem 1, the weakest simplicity standard in our class is obvious dominance of 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 in this section. If a game Γ admits a profile of obviously dominant strategies, then the game and the resulting mechanism are said to be **obviously strategy-proof (OSP)**.

In this section, we focus on environments without transfers and show that any OSP game is equivalent to a *millipede game*.[26] Roughly speaking, a millipede game is a clinch-or-pass game similar to a centipede game (Rosenthal, 1981), but with possibly more players and more actions ("legs") at each node. A simple example of a millipede game in an object allocation environment is a **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. Section 4.3).

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 $\bar{h} \supseteq h$ at which 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 **clinchable** 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 is notationally convenient 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 $P_i(h)$, $C_i^{\subseteq}(h)$ and $C_i^{\subset}(h)$ 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 $h' \subsetneq h$ and either $i_h = i$ or h is a terminal history. We say that payoff x **becomes impossible** for i at h if $x \in P_i(h')$ for all $h' \subsetneq h$ such that $i_{h'} = i$, but $x \notin P_i(h)$. We say payoff x is **previously unclinchable** at h if $x \notin C_i^{\subset}(h)$.

[26] This characterization complements Li's (2017b) result that in binary allocation environments with transfers, every OSP mechanism is equivalent to a personal clock auction; cf. Section 4.2 below.

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.[27]

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 $h_\varnothing$, 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 unclinchable payoff x that becomes impossible for agent i_h at h, then $C^{\complement}_{i_h}(h) \subseteq C_{i_h}(h)$.

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:

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.

[27]A 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.

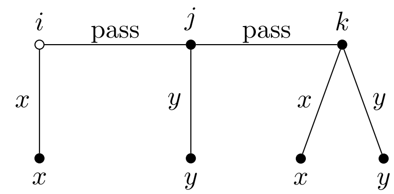

Figure 1: An example of a non-dictatorial millipede game in a voting environment with two outcomes, $\mathcal{X} = \{x, y\}$. The obviously dominant (greedy) strategy profile is for any agent to clinch if she is offered to clinch her preferred option among $\{x, y\}$, and otherwise pass.

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.[28] 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" 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. 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.

[28]Both parts of this transformation were first asserted for OSP in a footnote by Ashlagi and Gonczarowski (2018); cf. our Theorem 4 and Lemma A.4.

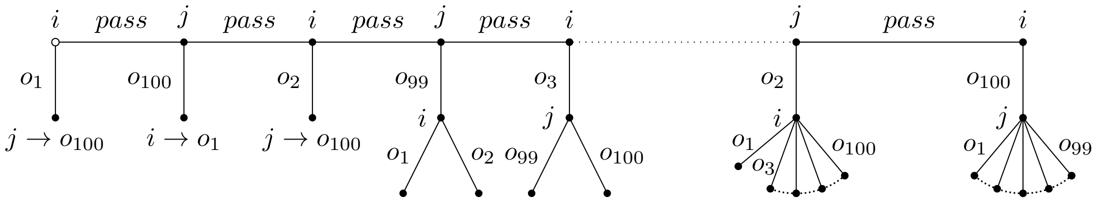

Figure 2: An example of a millipede game with two agents $\{i, j\}$ and 100 objects $\{o_1, o_2, \ldots, o_{100}\}$. If the first clinching is in an agent's first 50 moves, then the other agent is given the choice of clinching any object he or she could have clinched previously; if the first clinching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

4.2 One-Step Dominance

One-step simple dominance is stronger than obvious dominance. To see why this strengthening 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 $o_{100} \succ_i o_1 \succ_i o_2 \succ_i \cdots \succ_i o_{99}$ 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.[29]

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

Consider a set of outcomes $\mathcal{X} = Y \times \mathbb{R}^{\mathcal{N}}$, 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 $w = (w_i)_{i \in \mathcal{N}}$. Agents have types $\theta_i \in [\underline{\theta}_i, \overline{\theta}_i]$, where $0 \le \underline{\theta}_i < \overline{\theta}_i < \infty$,

[29]The first 100 moves of this millipede cannot be substantially shortened because, given the players' greedy strategies, for $k = 1, \ldots, 50$, 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.

and each agent's preferences are represented by a quasilinear utility function: $u_i(\theta_i, y, w) = \theta_i y_i + w_i$. Following Li (2017b), we call this preference environment binary allocation with transfers.[30] 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 Supplementary Appendix B.3, where we also prove Theorem 6.

Li (2017b) shows that in binary allocation settings, OSP games are equivalent to personal clock auctions. We strengthen this result to show that personal clock auctions are also OSS. Thus, in the binary setting, 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 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 by generalizing the one-step dominant strategic plans from the ascending auction example of Section 3. As in the ascending auction, in a personal clock auction, whenever an agent's price changes, she is offered an opportunity to drop out. In effect, the strategic plan to stay in whenever an agent's personal price is below her valuation and to plan to drop out at any next information set is one-step dominant.

[30] We 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.

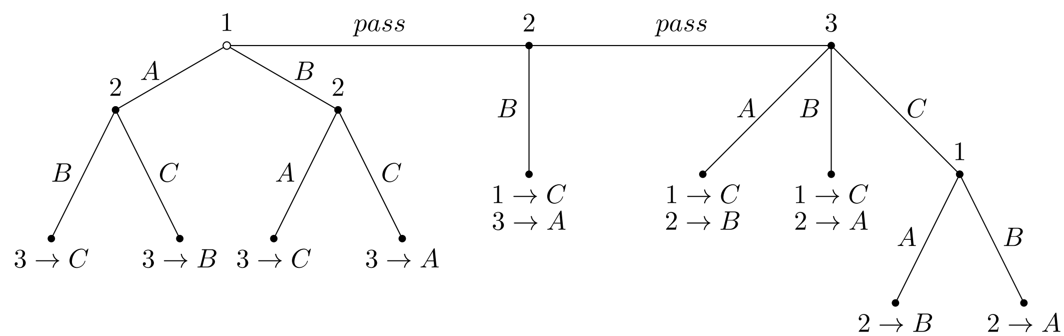

Figure 3: An example of a monotonic millipede game. The game allocates three objects A, B, and C to three agents (or players) $1, 2$ and 3. Agent 1 moves first and can clinch one of the objects A or B, or can pass. The second move is made by agent 2, who either clinches an object (in which case the allocation 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.

Environments without transfers

Recall the complex millipede game of Figure 2 that requires lengthy foresight on the part of the agents. One-step simplicity eliminates these complex millipede games, and leaves only millipedes 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 $h \subsetneq h'' \subsetneq h'$, either (i) $C_i(h) \subseteq C_i(h')$ or (ii) $P_i(h) \setminus C_i(h) \subseteq C_i(h')$. In words, this says that 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 that allows for passing actions and more complex allocation rules, see Figure 3. 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 with one-step dominant strategic collections. 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 outcome that is possible 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 outcome 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 everything possible but currently unclinchable. 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 the agent is 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,[31] but ascending auctions are not. Thus, SOSP further delineates the class of games that are simple to play, by eliminating millipede games that require even one-step forward-looking

[31] By Random Priority we mean the following mechanism that is commonly used to allocate indivisible objects to a group of agents: first, Nature randomly selecting an ordering of the agents, and then the agents are called one at a time in this order to select their favorite still-remaining object at their turn (cf. Abdulkadiroğlu and Sönmez (1998), Bogomolnaia and Moulin (2001), and Liu and Pycia (2011) who study the direct mechanism implementation of Random Priority). This mechanism is also sometimes called Random Serial Dictatorship, and is a special case of the more general sequential choice mechanisms we characterize in this section.

behavior. As there is a one-to-one correspondence between strongly obviously dominant strategic collections and strongly obviously dominant strategies, for simplicity of exposition we focus on strategies. Additionally, in this section we make use of the concept of an undominated payoff, where we say that a payoff x is **undominated** in a subset of payoffs for agent i if there is no payoff y in this subset such that $y \rhd_i x$. A mechanism $(\Gamma, (S_i(\succ_i))_{i \in \mathcal{N}})$ is **pruned** if every information set in Γ is on the path of play for some type of some player. Li (2017b) observed that every OSP mechanism is equivalent to a pruned OSP mechanism. The same is true for SOSP, cf. Appendix A.1.

Strongly obviously strategy-proof games are particularly simple to play. Any strongly obviously dominant strategy is greedy. Further, SOSP games can be implemented so that each agent is called to move at most once and has 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 game path of a pruned SOSP mechanism, there is at most one payoff-relevant history for each agent.*

This result—proven in Supplementary Appendix B.5—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. While an agent might be called to act later in the game, and her choice might influence the continuation game and the payoffs for other agents, it cannot affect her own payoff.

Building on Lemma 1, we show that SOSP effectively implies that agents—in a sequence—are faced with choices from personalized menus, e.g., in allocation with transfers this may be menus of object-price pairs. At the typical payoff-relevant history an agent is offered a menu of payoffs that she can clinch, she selects one of the alternatives from the menu, and she is never called to move again. More formally, we say that Γ is a **sequential 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.

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 priority mechanism (or, 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 priority mechanisms is that at an agent's turn in sequential choice, 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 prices, 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.[32]

[32] Prior 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.

5 Conclusion

We study the question of what makes a game simple to play, and introduce a general class of simplicity standards that vary the planning horizons of agents in extensive-form imperfect-information games. We allow agents that form a strategic plan only for a limited horizon in the continuation game, and the agents may update these plans as the game progresses and the future becomes the present. The least restrictive simplicity standard included 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. We show that SOSP delineates a class of posted-price mechanisms, OSS delineates a class of ascending clock mechanisms, and OSP delineates a richer class of mechanisms we call millipedes.[33] 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).[34] At the same time, strong obvious dominance is more restrictive, and more severely limits the class of implementable objectives. In environments with transfers, SOSP generally precludes efficiency and revenue maximization.[35] In environments without transfers, however, even SOSP

[33]Building on these results, Pycia and Troyan (2023) show that every obviously strategy-proof, Pareto efficient, and symmetric mechanism is equivalent to Random Priority.

[34]In Li's binary allocation settings, we show that all OSP mechanisms can be simplified to OSS.

[35]For instance, when we want to allocate an object to the highest value agent with transfers and with 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. This also 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.

mechanisms—serial dictatorships—can achieve efficient outcomes. Building on the results of the present paper, in single-unit demand allocation problems without transfers, Pycia and Troyan (2023) and Pycia (2017) show that 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, imposing more restrictive simplicity standards on the mechanisms limits the set of implementable social choice functions.[36]

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 and Chakraborty and Kendall, 2022),[37] 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 future 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.

[36]A different approach to the trade-off between simplicity and flexibility was proposed by Li and Dworczak (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.

[37]For 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 (2019): "All Sequential Allotment Rules Are Obviously Strategy-Proof," BGSE Working Paper 1108.

——— (2020): "On obvious strategy-proofness and single-peakedness," *Journal of Economic Theory*, 186, 1–34.

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.

Bewley, T. F. (1987): "Knightian Decision Theory, Part II: Intertemporal Problems," Cowles Foundation Discussion Paper 835.

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.

CATONINI, E. AND J. XUE (2021): “Local Dominance,” Work in Progress.

CHAKRABORTY, A. AND C. KENDALL (2022): “Future Self-Proof Elicitation Mechanisms,” Working Paper.

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): “(Il) 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.

Glazer, J. and A. Rubinstein (1996): "An Extensive Game as a Guide for Solving a Normal Game," *Journal of Economic Theory*, 70, 32–42.

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.

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.

Knight, F. H. (1921): *Risk, Uncertainty, and Profit*, Houghton Mifflin Company, Boston, MA.

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. (2020): "A Revelation Principle for Obviously Strategy-proof Implementation," *Games and Economic Behavior*, 124, 512–533.

Mandal, P. and S. Roy (2020): "Obviously Strategy-proof Implementation of Assignment Rules: A New Characterization," .

——— (2021): "On Obviously Strategy-proof Implementation of Fixed Priority Top Trading Cycles with Outside Options," *Economics Letters*, Forthcoming.

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. (2015): "Making Just School Assignments," *Games and Economic Behavior*, 92, 18–27.

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 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 P. Troyan (2023): "The Random Priority Mechanism is Uniquely Simple, Efficient, and Fair," 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.

Root, J. and D. S. Ahn (2020): "Incentives and Efficiency in Constrained Allocation Mechanisms," Working Paper.

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

Roth, A. 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.

Strotz, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 165–180.

Troyan, P. (2019): "Obviously Strategy-Proof Implemenation 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," Working Paper.

A Appendix: Proofs

This appendix contains the central elements of the proofs of our main theorems. All lemmas used in these proofs, as well as Theorem 6 and Lemma 1 from the main text, are proven in the Supplementary Appendix.

A.1 Pruning Principle

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. Li (2017b) introduced the following **pruning principle**: if $(S_i(\succ_i))_{i\in\mathcal{N}}$ is obviously dominant for Γ, then 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: if $(S_i(\succ_i))_{i\in\mathcal{N}}$ is strongly obviously dominant for Γ, then the restriction of $(S_i(\succ_i))_{i\in\mathcal{N}}$ to its pruning Γ' is strongly obviously dominant for Γ', and both games result in the same outcome.

A.2 Proof of Theorem 2

In light of Corollary 1, it is sufficient to prove the result for $k < \infty$ and $k' = k + 1$. For $k = 0$, the result 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, SOSP ($k = 0$) mechanism does not exist when there are at least two agents whose valuations are drawn iid from at least 3 values (see also footnote 35). For $k = 1$ we construct below a 2-step simple social choice rule that cannot be one-step implemented; we conclude the proof by extending this example to any larger k.

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 ($x, \tilde{x}$, etc.). The notation such as "$\ell \to \gamma$" below terminal nodes denotes that agent ℓ is assigned to object γ at this node, without needing to take any action. The root of the game is agent i's choice between clinching x and passing. 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.

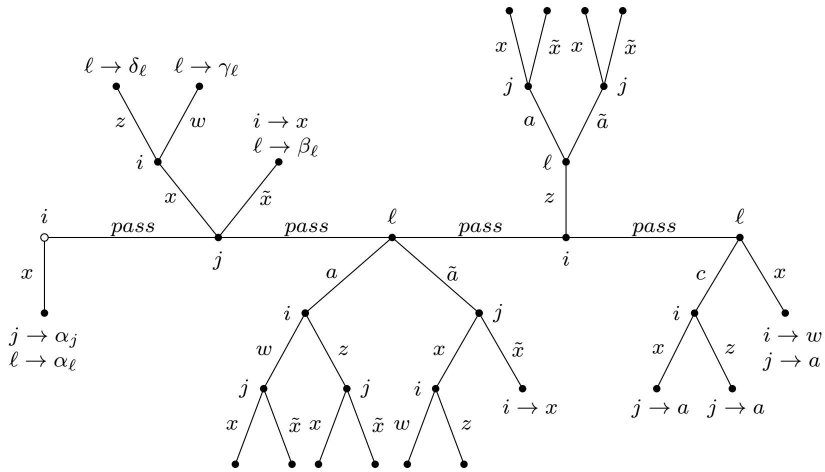

Figure 4: A game in which greedy strategies are two-step simple and for which no equivalent one-step simple mechanism exists.

It is straightforward to check that $(\Gamma, S_{\mathcal{N},\mathcal{H}})$, where $S_{\mathcal{N},\mathcal{H}}$ is a profile of greedy strategic collections, is k-step implementable for any $k \geq 2$; in particular, this implies that Γ is OSP. It is also easy to check that Γ itself is not OSS: the type of i that ranks $w \succ_i x \succ_i z$ has no one-step simple strategic plan when choosing between x and passing at the first move of the game. Showing that the social choice rule implemented by $(\Gamma, S_{\mathcal{N},\mathcal{H}})$ cannot be OSS-implemented by any other mechanism is subtler, and we relegate the proof of the following lemma establishing this statement to Supplementary Appendix B.1.

Lemma A.1. *No one-step simple mechanism is equivalent to* $(\Gamma, S_{\mathcal{N},\mathcal{H}})$.

For $k = 2$, game $\Gamma^{(2)}$ in Figure 5 is an example that is k'-step simple for any $k' > k$, but for which no equivalent k-step simple mechanism exists. This game is similar in structure to that of Figure 4, but has the following additions:

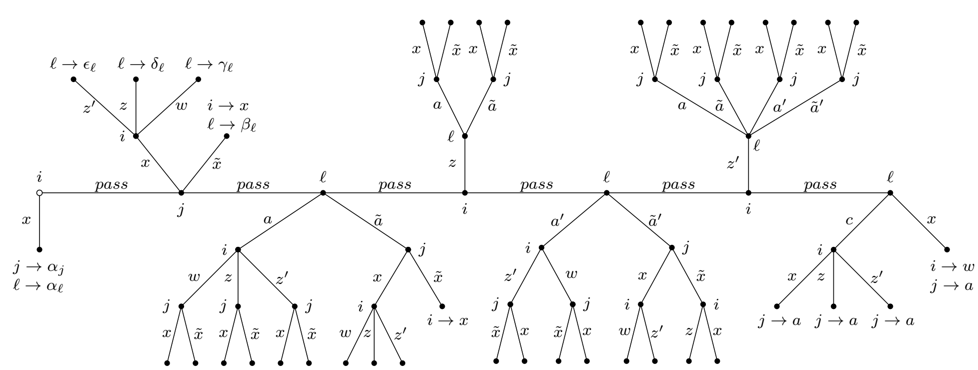

Figure 5: A game in which greedy strategies are three-step simple and for which no equivalent two-step simple mechanism exists.

(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.[38]

(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 prove the theorem for arbitrary $k \geq 2$, 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.

It is straightforward to check that $\left(\Gamma^{(k)}, S^{(k)}_{\mathcal{N},\mathcal{H}}\right)$ is $(k+1)$-step simple but not k-step simple, where $S^{(k)}_{\mathcal{N},\mathcal{H}}$ is a profile of greedy strategic collections. Showing that no equivalent mechanism is k-step simple is done similarly to the $k = 1$ case. The details can be found in Supplementary Appendix B.1.

Lemma A.2. *For any $k \geq 2$, no k-step-simple mechanism is equivalent to $\left(\Gamma^{(k)}, S^{(k)}_{\mathcal{N},\mathcal{H}}\right)$.*

Lemmas A.1 and A.2 establish the result for $k \geq 1$. ■

A.3 Proof of Theorem 3

The proof develops 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 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

[38] This property and the property 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 possible for i.

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 $\gtrsim_i$ 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 $\gtrsim_i$ -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 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

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 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.

The proof is broken down into five steps, stated as Lemmas A.3-A.7 below. The proofs of these lemmas can be found in Supplementary Appendix B.2. First, we check there that all millipede games with greedy strategies are OSP, establishing one direction of the theorem.

Lemma A.3. *Millipede games with greedy strategies are obviously strategy-proof.*

Given Li's pruning principle (see Subsection A.1), the converse implication of Theorem 5—that all OSP mechanisms are equivalent to millipedes—follows from the remaining four lemmas.[39] Lemma A.4 develops Theorem 4 (see this theorem for a discussion):

Lemma A.4. *Every OSP game is equivalent to an OSP game with perfect information in which Nature moves at most once, as the first mover.*

Lemma A.5 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.[40]

Lemma A.5. *Let Γ be an obviously strategy-proof game of perfect information that is pruned with respect to the obviously dominant strategy profile $(S_i(\succ_i))_{i\in\mathcal{N}}$. Consider a history h where agent $i_h = i$ is called to move. There is at most one action $a^* \in A(h)$ such that $P_i((h, a^*)) \notin G_i(h)$.*

The above lemma 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.6. *For any OSP game Γ, there exists an equivalent OSP game Γ' such that the following hold at each $h \in \mathcal{H}$ (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 $a \in A(h)$ are clinching actions and $i_{h'} \neq i$ for any $h' \supsetneq h$.

The final lemma of the proof establishes the payoff guarantees in the game constructed in the previous lemmas.

Lemma A.7. *Let $(\Gamma, S_{\mathcal{N}})$ be an obviously strategy-proof mechanism that satisfies the conclusions of Lemmas A.4 and A.6. At all h, if there exists a previously unclinchable payoff z that becomes impossible for agent i_h at h, then $C^c_{i_h}(h) \subseteq C_i(h)$.*

This lemma concludes the proof of Theorem 5. ∎

[39] We 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.6 below).

[40] We emphasize the distinction between a payoff x being "guaranteeable" vs. "clinchable": the latter means the agent receives x at all terminal histories, while the former means there is a continuation strategy S_i such that she receives x at all terminal histories consistent with S_i.

A.5 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 $\bar{x}_{h^*} = Top(\succ_i, P_i(h^*))$ and $\bar{y}_{h^*} = Top(\succ_i, C_i(h^*))$. Let $\mathcal{H}_{i,h^*} = \{h \in \mathcal{H}_i | h^* \subsetneq h' \subsetneq h \implies h' \notin \mathcal{H}_i\}$ 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 $S_{i,h^*}(h^*) = a_{\bar{x}_{h^*}}$.
  - Else, we have $C_i(h^*) \subseteq C_i(h)$ (by monotonicity) and we set $S_{i,h^*}(h^*) = a_{\bar{y}_{h^*}}$.

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.

In order to 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.6. 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.6. It remains to show:

Lemma A.8. *Any OSS millipede mechanism that, at each $h \in \mathcal{H}$, satisfies conditions (i), (ii), and (iii) of Lemma A.6 is monotonic.*

We prove this lemma in Supplementary Appendix B.4. ∎

A.6 Proof of Theorem 8

That sequential choice mechanisms are SOSP is immediate from the definition, and so we focus on proving that every SOSP mechanism is equivalent to a sequential choice mechanism. Following the same reasoning as in the proof of Lemma A.4, given any SOSP mechanism, we can construct an equivalent SOSP mechanism of perfect information in which Nature moves at most once, as the first mover. It remains to analyze the subgame after a potential move by Nature and to show that every perfect-information SOSP mechanism in which there are no moves by Nature is equivalent to a sequential choice mechanism.

Let $(\Gamma, S_{\mathcal{N}})$ be such a mechanism. In line with the discussion in Section A.1, we can assume that Γ is pruned. By Lemma 1, each agent i can have at most one payoff-relevant history along any path of 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 Lemma A.6, we construct an equivalent pruned game 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. 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 construct a pruned 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. It remains to show

Lemma A.9. Γ' *with greedy strategies is a sequential choice mechanism.*

We prove this lemma in Supplementary Appendix B.6. ■

B Supplementary Appendix: Omitted Proofs (For Online Publication)

This supplementary appendix contains the proofs of the lemmas used in the proofs of the main theorems in Appendix A, as well as the full proofs of Theorem 6 and Lemma 1 from the main text.

B.1 Proofs of Lemmas for Theorem 2

Proof of Lemma A.1. In order to show that there is no OSS mechanism that is equivalent to Γ, 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 use analogues of these steps to prove Lemma A.2 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.[41] 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 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

[41] That 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).

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.[42] 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 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

[42] Since 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.

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.[43] *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 no obviously dominant action at ℓ's prior moves, which is a contradiction.[44]

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

[43] The argument shows that x not only is not clinchable for i but also not guaranteeable.

[44] Note 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.

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. ∎

Proof of Lemma A.2. Take any $k \geq 2$. By way of contradiction, 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 begins by repeating steps 1.1-1.6 from the proof of Lemma A.1 above, 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 $k' = 2, 3, \ldots, k$. 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. ∎

B.2 Proofs of Lemmas for Theorem 5

Proof of Lemma A.3. 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 $x_W(h) \succsim_i x_W(h')$ for all $h' \subsetneq h$ such that $i_{h'} = i$. Since passing is obviously dominant at all $h' \subsetneq h$, we have $x_W(h') \succsim_i x_C(h')$, and together, these imply that $x_W(h) \succsim_i x_C(h')$ for all 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 $h' \subsetneq h$ such that $i_{h'} = i$.

Let $X_0 = \{x' : x' \in P_i(h) \text{ and } x' \succ_i x_C(h)\}$; in words, X_0 is a set of payoffs that are possible at all $h' \subseteq h$, and are strictly better than anything that was clinchable at any $h' \subseteq h$ (and therefore have never been clinchable themselves). Order the elements in X_0 according to $\succ_i$, and without loss of generality, let $x_1 \succ_i x_2 \succ_i \cdots \succ_i x_M$.

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) $i_h = i$ 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 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 $m' = 1, \ldots, m-1$, 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 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): $h' \subsetneq \hat{h}_{-m}$. 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 $\{x_1, \ldots, x_{m-1}\}$ are previously unclinchable at $\hat{h}_{-m}$, and so $x_{m'}$ is previously unclinchable at $\hat{h}_{-m}$. By definition of a millipede game (part 3), we 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 of Lemma A.3, let $\hat{h} = \max\{h_1, h_2, \ldots, h_M\}$ (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 unclinchable at

[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 $x_W(h)$.

$\hat{h}$. 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). ∎

Proof of Lemma A.4. 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 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 $\varnothing$. 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 $\varnothing$ by Nature in Γ' as an action made by Nature at h in game Γ). When the support of possible outcomes shrinks, the worst-case outcome from any fixed strategy can only improve, while the best-case can only diminish, and so if a strategy was obviously dominant in Γ, the corresponding strategy will continue to be obviously dominant in Γ', and the two games will be equivalent. ∎

Proof of Lemma A.5. 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 $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \varnothing$, 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 $\succ_i$ is following some strategy such that $S_i(\succ_i)(h) = a_1^*$ that results in a payoff of x at some terminal history $\bar{h} \supseteq (h, a_1^*)$. Note that $Top(\succ_i, P_i(h)) \neq x$, because otherwise a_1^* would not be obviously dominant for this type (since $x \notin G_i(h)$ and $x \in P_i(h_2^*)$). Thus, let $Top(\succ_i, P_i(h)) = y$. Note that $y \notin G_i(h)$ (or else it would not be obviously dominant for type $\succ_i$ to play a strategy such that x is a possible payoff). 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 $\succ_i$, which contradicts that $S_i(\succ_i)(h) = a_1^*$; given the former, if $y \in P_i(h_2^*)$, then once again a_1^* would not be obviously dominant for type $\succ_i$. Now, again by pruning, there must be some type $\succ_i'$ such that $S_i(\succ_i')(h) = a_2^*$ that results in payoff x at some terminal history $\bar{h} \supseteq (h, a_2^*)$. By similar reasoning as previously, $Top(\succ_i', P_i(h)) \neq x$, and so $Top(\succ_i', P_i(h)) = z$ for some $z \in P_i(h_2^*)$. Since $y \notin P_i(h_2^*)$, we have $z \neq y$, and we can as above conclude that $z \notin G_i(h)$. It is without loss of generality to consider a type $\succ_i'$ such that $Top(\succ_i', P_i(h) \setminus \{z\}) = y$. 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, $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \varnothing$, 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 $Top(\succ_i, P_i(h)) \neq x, y$. To see why, assume there were a type that reaches h such that $Top(\succ_i, P_i(h)) = x$. Then, by richness, there is a type that reaches h such that $Top(\succ_i, P_i(h) \setminus \{x\}) = y$. 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 $S_i(\succ_i)(h) = a_1^*$ and x is a possible payoff. Let $Top(\succ_i, P_i(h)) = z$ for this type, where, as just noted, $z \neq x, y$. The fact that $S_i(\succ_i)(h) = a_1^*$ implies that $z \in P_i(h_1^*)$ and $z \notin G_i(h)$; if either of these were false, it would not be obviously dominant for this type to play a strategy such that $S_i(\succ_i)(h) = a_1^*$ and x is a possible payoff. In other words, $z \in PnG(h)$ and $z \in P_i(h_1^*)$. Since we just showed that $PnG_i(h) \cap P_i(h_1^*) \cap P_i(h_2^*) = \varnothing$, we have $z \notin P_i(h_2^*)$. Finally, consider a type $\succ_i$ such that $Top(\succ_i, P_i(h)) = z$ and $Top(\succ_i, P_i(h) \setminus \{z\}) = y$. Note that this type has no obviously dominant action at h, which is a contradiction. ■

Proof of Lemma A.6. Given an OSP mechanism $(\Gamma, S_{\mathcal{N}})$, begin by using Lemma A.4 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 $S_{\mathcal{N}}$. With slight abuse of notation, we denote this pruned, perfect information mechanism by $(\Gamma, S_{\mathcal{N}})$. Consider some history h of Γ at which the mover is $i_h = i$. By Lemma A.5, all but at most one action (denoted a^*) in $A(h)$ satisfy $P_i((h, a)) \subseteq G_i(h)$; this means that any obviously dominant strategy for type $\succ_i$ that does not choose a^* guarantees the best possible 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 $a_1, \ldots, a_M$. 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.[46]

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 agents follow strategies in the new game analogous to the their strategies from the original game, the same outcome is reached, and so Γ and Γ' are equivalent under their respective strategy profiles.

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

[46] More 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.

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. ■

Proof of Lemma A.7. Let h be any earliest history where some agent i moves such that there is a previously unclinchable payoff z that becomes impossible at h (the case for terminal histories is 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 $i_h = i$; and (c) $z \notin C_i^{\subset}(h)$. 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; thus, for any type of agent i that ranks z first, any obviously dominant strategy has the agent choosing the unique passing action at all $h' \subsetneq h$.

We want to show that $C_i^{\subset}(h) \subseteq C_i(h)$. Towards a contradiction, assume that $C_i^{\subset}(h) \nsubseteq C_i(h)$, and let $x \in C_i^{\subset}(h)$ 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.6, 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 $h' \subsetneq \bar{h}$ 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 $h'' \subsetneq \bar{h}$ 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. However, at h', i could have clinched x, and so the strategy is not obviously dominant, a contradiction. ■

B.3 Proof of Theorem 6

Before proving the theorem, we first formally define 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.[47]

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 $\tilde{w}_i : \{h_i : h_i^* \subseteq h_i\} \to \mathbb{R}$ 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 $g_{w,i}(\bar{h}) = \inf\{\tilde{w}_i(h_i) : h_i^* \subseteq h_i \subsetneq \bar{h}\}$.
- If $\bar{h} \supsetneq (h, a)$ for some $h \in \mathcal{H}_i$ and $a \in A^q$, then $i \notin g_y(\bar{h})$.
- $A^q \cap A(h_i^*) \neq \varnothing$
- For all $h_i', h_i'' \in \{h_i \in \mathcal{H}_i : h_i^* \subseteq h_i\}$:
  - If $h_i' \subsetneq h_i''$, then $\tilde{w}_i(h_i') \geq \tilde{w}_i(h_i'')$
  - If $h_i' \subsetneq h_i''$, $\tilde{w}_i(h_i') > \tilde{w}_i(h_i'')$ and there is no h_i''' such that $h_i' \subsetneq h_i''' \subsetneq h_i''$, then $A^q \cap A(h_i'') \neq \varnothing$
  - If $h_i' \subsetneq h_i''$ and $\tilde{w}_i(h_i') > \tilde{w}_i(h_i'')$, then $|A(h_i') \setminus A^q| = 1$
  - If $|A(h_i') \setminus A^q| > 1$, then there exists $a \in A(h_i')$ such that, for all $\bar{h} \supseteq (h_i', a)$, $i \in g_y(\bar{h})$;[48]

or, **Out-Transfer Falls**: as above replacing 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, Li (2017), and the construction of the one-step simple strategic

[47] We slightly simplify Definition 15 of Li (2017b) by restricting it to perfect information games: by Theorem 4, for any personal clock auction that satisfies Definition 15 of Li (2017b), there is an equivalent mechanism that satisfies the definition we work with. This also applies to the minor correction provided by Li in a corrigendum available on his website; cf. footnote 48 for further details.

[48] The 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.

collections 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 construction and argument for in-transfer falls; the case of out-transfer falls is analogous. Consider any $h_i \in \mathcal{H}_i$ and simple-node set $\mathcal{H}_{i,h_i} = \{h' \in \mathcal{H}_i : h_i \subsetneq h'' \subsetneq h' \implies h'' \notin \mathcal{H}_i\}$, and define the strategic plan $S_{i,h_i}(h')$ at $h' \in \mathcal{H}_{i,h_i}$ as follows:

- If $\theta_i + \tilde{w}_i(h_i) > \bar{w}_i$ and $A(h_i) \setminus A^q \neq \varnothing$:
  - [**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 $a \in A(h') \cap A^q$. Else, $S_{i,h_i}(h') = a'$ for some $a' \in A(h')$ such that for all $\bar{h} \supseteq (h, a')$, $i \in g_y(\bar{h})$.
- Else, 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, first consider h_i such that $A(h_i) \setminus A^q = \varnothing$. Then the only actions available at h_i are quitting actions. Thus, the best- and worst-cases from any action are all $\bar{w}_i$, and one-step dominance holds. Second, consider $\theta_i + \tilde{w}_i(h_i) \leq \bar{w}_i$. 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. Third, consider the remaining case in which $\theta_i + \tilde{w}_i(h_i) > \bar{w}_i$ and there exists some $a \in A(h_i) \setminus A^q$. There are two subcases:

First, if $|A(h_i) \setminus A^q| = 1$, then all other 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., $A(h'_i) \cap A^q \neq \varnothing$), 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., $A(h'_i) \cap A^q = \varnothing$), 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.

Second, 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. ∎

B.4 Proof of Lemma for Theorem 7

Proof of Lemma A.8. By way of contradiction, let $(\Gamma, S_{\mathcal{N},\mathcal{H}})$ be a millipede mechanism that satisfies (i)-(iii) at each history but is not monotonic. The failure of monotonicity implies that there exists an agent i, history $h^* \in \mathcal{H}_i$, history h that follows i's passing move at h^* that is either terminal or in $\mathcal{H}_i$ and such that i does not move between h^* and h, and payoffs x and y such that $x \in (P_i(h^*) \smallsetminus C_i(h^*)) \smallsetminus C_i(h)$ and $y \in C_i(h^*) \smallsetminus C_i(h)$; in particular, $x \neq y$. Without loss of generality, assume that h^* is an earliest history at which monotonicity is violated in this way. This implies that $x \notin C_i(h')$ for any $h' \subseteq h^*$ such that $i_{h'} = i$.[49] In particular, any type $\succ_i$ of agent i that ranks payoff x first passes at any $h' \subseteq h^*$ at which this agent moves.

As $x, y \notin C_i(h)$ by the choice of these payoffs, there is some third payoff $z \neq x, y$ such that $z \in C_i(h)$. Let $\succ_i$ be such that $\succ_i: x, y, z \ldots$ and $\succ_i'$ be such that $\succ_i': x, z, \ldots$; these types exist by richness, given that we are in a no-transfer environment. Ranking x first, these types are passing at all nodes $h' \subseteq h^*$ at which they move: $S_{i,h'}(\succ_i)(h') = S_{i,h'}(\succ_i')(h') = a^*(h')$ where $a^*(h')$ denotes the passing action at h'.

We conclude the indirect argument by showing that none of the following two cases is possible:

Case $y \notin P_i(h)$. If also $x \notin P_i(h)$, then $P_i(h)$ contains some $w \neq x, y$. If $x \in P_i(h)$, then $x \notin C_i(h)$ implies that $x \in P_i((h, a^*(h)))$ and by definition of a passing action, there is some $w \neq x$ such that $w \in P_i((h, a^*))$; furthermore $w \neq y$ because $y \notin P_i(h)$. In either case, passing at h^* might lead to w which is worse for $\succ_i$ than y, and i can clinch y at h^*; thus $S_{i,h^*}(\succ_i)$, which passes at h^*, is not one-step dominant; a contradiction.

Case $y \in P_i(h)$. If $z \in P_i((h, a^*))$ then $x, y \notin C_i(h)$ implies that the worst case for type $\succ_i$ from passing at h^* is at best z, which is worse than clinching y at h^*. Therefore, the passing action $S_{i,h^*}(\succ_i)$ is not one-step dominant at h^* for $\succ_i$, a contradiction. We may thus assume that $z \notin P_i((h, a^*))$. Because $x \notin C_i(h)$, the assumptions of Lemma A.6 imply that x is not guaranteeable at h, and in particular it is not guaranteeable at $(h, a^*(h))$. Thus,

[49] If $x \in C_i(h')$ for some h', then, by monotonicity, at any next history $h'' \supsetneq h'$ following a pass where i moves, either $x \in C_i(h'')$ or $P_i(h') \smallsetminus C_i(h') \subseteq C_i(h'')$. If the latter holds, then at h'', i has been offered to clinch everything that is possible for her, and so, by greediness, 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.

the worst case for type $\succ_i'$ from passing at h is strictly worse than z; since $z \in C_i(h)$, this implies that $S_{i,h}(\succ_i')$ clinches at h. Thus $x \notin C_i(h)$ allows us to conclude that $x \notin P_i(h)$, as otherwise $S_{i,h}(\succ_i')$ could not be clinching at h. Since $y \notin C_i(h)$ and $y \in P_i(h)$, we infer that $y \in P_i((h, a^*(h)))$. As at least two payoffs are possible following passing and $x \notin P_i(h)$, there is some $w \neq x, y$ that is possible at $(h, a^*(h))$ and hence also at h. As x is not possible and y is not clinchable at h, the worst case for type $\succ_i$ from the perspective of node h^* from following $S_{i,h^*}(\succ_i)$ is at best w, which is strictly worse than clinching y at h^*. Thus $S_{i,h^*}(\succ_i)$ is not one-step dominant. ■

B.5 Proof of Lemma 1

Recall that any strongly obviously dominant strategy is greedy. We first note the following lemmas. To state the lemmas, define $\hat{P}_i(h) = \{x \in P_i(h) : \nexists y \in P_i(h) \text{ s.t. } y \rhd_i x\}$ to be the set of possible payoffs for i at h that are undominated in $P_i(h)$.

Lemma B.1. *Let Γ be a pruned SOSP game. If a history h at which agent i moves is payoff-relevant, then $|\hat{P}_i(h)| \geq 2$.*

Proof of Lemma B.1. Assume not, and let $\hat{P}_i(h) = \{x\}$, where x is the unique undominated payoff at h.[50] In particular, $x \rhd_i x'$ for all $x' \in P_i(h)$, and $Top(\succ_i, P_i(h)) = x$ for all types of agent i. Because x is possible at h, there is an action $a \in A(h)$ such that $x \in P_i((h, a))$. Action a does not clinch x; indeed if $P_i((h, a)) = \{x\}$ then greediness would imply that only actions clinching x are taken, and in a pruned game h would not be payoff relevant. Thus, there is another $x' \in P_i((h, a))$ such that $x \succ_i x'$ for all types of agent i. Let $a' \neq a$ be an action at $A(h)$. If $x \in P_i((h, a'))$, then, analogously as for a, there is some other $x'' \in P_i((h, a'))$. It is then easy to check that neither a nor a' strongly obviously dominates the other. If $x \notin P_i((h, a'))$ then it would not be strongly obviously dominant (SOD, for shortness) for any type to take action a', which would contradict the game being pruned. ■

Lemma B.2. *Let (Γ, S) be a pruned SOSP mechanism. 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 B.2. 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 $a \in A(h_0^i)$. By way of contradiction assume that there exists some 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. As (Γ, S) is pruned, there must

[50] There must be at least one undominated payoff, since $\trianglerighteq_i$ is transitive and the number of payoffs is finite.

be some type, $\succ_i$, for which action a is SOD at h_0^i. Possibly by renaming the outcomes, richness allows us to assume that $Top(\succ_i, P_i(h_0^i)) = x$ and $x \succ_i y \succ_i z$. For a to be strongly obviously dominant, for all other actions $a' \neq a$ at h_0, the best case outcome for type $\succ_i$ following a' must be no better than z; in particular, this implies that for all $a' \neq a$ and all $w \in \hat{P}_i((h_0^i, a'))$, $w \ntrianglerighteq_i y$. Let $a'' \neq a$ be an action at h_0. If there is $w \in \hat{P}_i((h_0^i, a''))$ such that $x \ntrianglerighteq_i w$, then there is a type $\succ_i'$ such that $Top(\succ_i', P_i(h_0^i)) = y$ and $y \succ_i' w \succ_i' x$. For this type, the worst case from a is at best x, while w is possible following a'' , so a is not strongly obviously dominant; for any $a' \neq a$, the worst case is strictly worse than y as argued above, while the best case from a is y, and so no $a' \neq a$ is SOD either. Therefore, type $\succ_i'$ has no SOD action, a contradiction showing that no $w \in \hat{P}_i((h_0^i, a''))$ satisfies $x \ntrianglerighteq_i w$. An analogous argument—with z playing the role of x—shows that no $w \in \hat{P}_i((h_0^i, a''))$ satisfies $z \ntrianglerighteq_i w$. Thus, for all a'' and all $w \in \hat{P}_i((h_0^i, a''))$, $x \trianglerighteq w$ and $z \trianglerighteq w$. As x and z are distinct, for any type $\succ_i'$, either $x \succ_i' w$ or $z \succ_i' w$, and in either case a'' is not a dominant action for a type contrary to (Γ, S) being pruned. This contradiction shows that $|\hat{P}_i((h_0^i, a))| \leq 2$ for all $a \in A(h_0^i)$.

Finally, we show that $|\hat{P}_i((h_0^i, a))| = 2$ for at most one $a \in A(h_0^i)$. Towards a contradiction, let a and a' be two actions such that there are two possible undominated payoffs for i following each, and, for notational purposes, let $\hat{P}_i((h_0^i, a)) = \{x, y\}$ and $\hat{P}_i((h_0^i, a')) = \{w, z\}$, where, a priori, it is possible that $w, z \in \{x, y\}$. As the mechanism is pruned, there is some type $\succ_i$ that selects action a as an SOD action; 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 SOD, the best case from any $a' \neq a$ must be no better than y, which implies that $w, z \ntrianglerighteq_i x$, and thus $x \neq w, z$. Pruning also implies that some type $\succ_i'$ is selecting action a'' as an SOD action; without loss of generality, let $Top(\succ_i', P_i(h_0^i)) = z$. Since w is possible following a'', in order for a'' to be SOD, the best case from a must be no better than w for type $\succ_i'$, thus $x, y \ntrianglerighteq_i z$, and so $z \neq x, y$. Thus, we have shown that x, y, z are all distinct, that no outcome in $P_i(h_0^i)$—including z and y—structurally dominates x, and that $y \ntrianglerighteq_i z$. Richness then implies that there is a type $\succ_i$ such that $Top(\succ_i, P_i(h_0^i)) = x$ and $x \succ_i z \succ_i y$. This type has no SOD action: only a can be SOD because only a makes x possible, but a is not SOD because the worst case from a is at best y, while the best case from a' is z. ■

Continuing with the proof of Lemma 1, assume that there was a path of the game with two payoff-relevant histories $h_1 \subsetneq h_2$ for some agent i. 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. First, we claim that there are at least two structurally undominated payoffs at h_1, i.e., $|\hat{P}_i(h_1)| \geq 2$. To show it by way of contradiction, suppose that $\hat{P}_i(h_1) = \{x\}$, which implies that $x \triangleright_i x'$ for all other $x' \in P_i(h_1)$. Then $P_i((h_1, a)) = \{x\}$ for all $a \in A(h_1)$. Indeed, suppose that

$x' \neq x$ is possible after some action a at h_1. Then $x, x' \in P_i((h_1, a))$ because otherwise no type of i finds a to be SOD, which is impossible as the game is pruned. If $x \in P_i((h_1, a'))$ for some action $a' \neq a$ at h_1, then a is not SOD for any type of i, which again is impossible as the game is pruned. Thus $x \notin P_i((h_1, a'))$ and no type of i finds a' to be SOD, which yet again is impossible in a pruned game. Thus, no $x' \neq x$ is possible after any $a \in A(h_1)$, which contradicts that h_1 is payoff-relevant. This contradiction shows that $\hat{P}_i(h_1)$, being non-empty, has at least two elements.

Let a_1^* be the action such that $h_2 \supseteq (h_1, a_1^*)$. By Lemma B.2, one of the below two cases would need to obtain, and to conclude the indirect argument we now show that neither of them obtains.

Case $|\hat{P}_i((h_1, a_1^*))| = 1$. Let z be the unique undominated payoff that is possible after a_1^*; $z \in \hat{P}_i(h_1)$ as otherwise no type of i would find a_1^* to be SOD, which is impossible in a pruned mechanism. Because h_2 is payoff-relevant, Lemma B.1 tells us that $|\hat{P}_i(h_2)| \geq 2$, and thus $z \notin \hat{P}_i(h_2)$ as z weakly structurally dominates all outcomes in $P_i(h_2) \subseteq P_i((h_1, a_1^*))$. Let $x \neq z$ be an outcome in $\hat{P}_i(h_1)$ and let $z', z'' \in \hat{P}_i(h_2)$ be distinct undominated payoffs that are possible at h_2, and consider a type $\succ_i$ that ranks the outcomes $z \succ_i x \succ_i z'$. For this type, a_1^* is not SOD at h_1 because z' is possible following a_1^* while $x \notin \{z\} = \hat{P}_i((h_1, a_1^*))$ is possible following some other action at h_1. No action $a \neq a_1^*$ is SOD for $\succ_i$ if $z \notin P_i((h_1, a))$. Hence $z \in P_i((h_1, a))$ but then a_1^* would not be SOD for any type; impossible as the mechanism is pruned. This contradiction shows that the present case is impossible.

Case $|\hat{P}_i((h_1, a_1^*))| = 2$. Then a_1^* is the unique action with two undominated payoffs from Lemma B.2; let us label these payoffs x and y. As the game is pruned, there is some type $\succ_i$ for which a_1^* is strongly obviously dominant; in particular the payoff $Top(\succ_i, P_i(h_1))$ is possible following a_1^* and by renaming payoffs we can set $x = Top(\succ_i, P_i(h_1))$. For each action $a \neq a_1^*$ at h_1, Lemma B.2 implies that $\hat{P}_i((h_1, a)) = \{w_a\}$, for some payoff w_a; action a_1^* being SOD for type $\succ_i$ implies that $w_a \not\succeq_i x$ (and in particular $w_a \neq x$); and a being SOD for some other type implies that $y \not\succ_i w_a$. If $w_a \neq y$ then $y \not\succeq_i w_a$, and, given that x and y are mutually undominated, richness would give us a type $\succ_i^a$ such that $x \succ_i^a w_a \succ_i^a y$, but for this type neither a_1^* nor a nor any other action a' at h_1 is SOD because as shown above $w_{a'} \neq x$. We conclude that $w_a = y$ for all actions $a \neq a_1^*$ at h_1.

To continue the indirect argument we now show that $\hat{P}_i(h_2) = \{x, y\}$. The set $\hat{P}_i(h_2)$ has two elements, by Lemma B.1, because h_2 is payoff relevant. Thus, if $\hat{P}_i(h_2) \neq \{x, y\}$ then there would be some $z \neq x, y$ such that $z \in \hat{P}_i(h_2) \subseteq P_i(h_1)$. As x and y are undominated at $(h_1, a_1^*) \subsetneq h_2$, richness would give us type $\succ_i^2$ such that $x \succ_i^2 y \succ_i^2 z$ and for this type a_1^* would not be SOD at h_1 because z would be possible following a_1^* while, as shown above, y would be possible following another action; further, no $a \neq a_1$ would be SOD at h_1 because

y would be possible following a while x would be possible following a_1^*. The lack of an SOD action is a contradiction showing that $\hat{P}_i(h_2) = \{x, y\}$. Thus any $z \in P_i(h_2)$ is structurally dominated by either x or y and for, each type, x or y is the top payoff in $P_i(h_2)$. Since $\hat{P}_i((h_1, a)) = y$ for all $a \neq a_1^*$, strong obvious dominance implies that all and only types $\succ_i^1$ with $x = Top(\succ_i^1, P_i(h_1))$ select action a_1^* at h_1 and hence these are the types for whom h_2 is on path. As y is possible at h_2, there is at least one action $a_2^* \in A(h_2)$ after which y is possible. As at each history agents have at least two actions, there is another action $a_2 \in A(h_2)$, and, as the mechanism is pruned, there are two types $\succ_i^{a_2^*}$ and $\succ_i^{a_2}$ for which h_2 is on path such that $\succ_i^{a_2^*}$ selects a_2^* and $\succ_i^{a_2}$ selects a_2 at h_2. Because we established that x is possible at h_2 and that it is the top possible payoff for both these types, SOSP implies that $x \in P_i((h_2, a_2^*))$ and $x \in P_i((h_2, a_2))$. By construction, $y \in P_i((h_2, a_2^*))$, and hence a_2^* is not SOD for type $\succ_i^{a_2^*}$; a contradiction that concludes the proof of the lemma. ■

B.6 Proof of Lemma for Theorem 8

Proof of Lemma A.9. By way of contradiction suppose that game Γ', together with greedy strategies, is not a sequential choice mechanism. Let h be an earliest history where the definition of a sequential choice mechanism is violated. As such h is payoff relevant and Γ' is pruned, Lemma 1 implies that h is a first history at which i moves. Since Γ' is not a sequential choice mechanism, there must be some payoff $x \in P_i(h)$ that i cannot clinch at h. We may assume that x is not dominated, i.e., $x \in \hat{P}_i(h)$; indeed, if all $x' \in \hat{P}_i(h)$ were clinchable at h, then greediness would imply that all dominated actions were pruned in Γ'. Since x is not clinchable, for any action $a \in A(h)$ such that $x \in P_i((h, a))$, there is some payoff in $P_i((h, a))$ that is different from x. We fix one such action a.

Case $|P_i(h)| = 2$. Let y be the other payoff in $P_i(h)$. If y were clinchable then the mechanism would satisfy the definition of sequential choice at h. Since we assumed that the definition is not satisfied at h, neither x nor y is clinchable. Thus, for all $a \in A(h)$, $P_i((h, a)) = \{x, y\}$. As x and y are different payoffs, at least one of $x \succ_i y$ or $y \succ_i x$ holds for some type at h. Because there are at least two actions in $A(h)$, this type does not have a strongly obviously dominant (SOD) action at h, which is a contradiction.

Case $|P_i(h)| \geq 3$ **and** $x \rhd_i y$ **for all** $y \neq x$ **in** $P_i((h, a))$**.** There is an action $a' \neq a$ at h and, because x is not clinchable at h, there is some $w \neq x$ that belongs to $P_i((h, a'))$. We have $y \unrhd_i w$; indeed, if not, then x being undominated implies that there would exist type $\succ_i$ such that $x \succ_i w \succ_i y$, and, taking into account that x is not clinchable at h, this type would have no SOD action at h. Thus, $x \rhd_i y \unrhd_i w$; but this implies that a' is not SOD for any type, which contradicts the mechanism being pruned.

Case $|P_i(h)| \geq 3$ and there exists $y \in P_i((h,a))$ such that x and y do not dominate each other. By Lemma B.2, for any $a' \neq a$, the set $\hat{P}_i((h,a'))$ is a singleton. 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 $\hat{P}_i((h,a')) = \{w'\}$. Then also $w' \neq x$; indeed, if $w' = x$ then, as x is both structurally undominated and unclinchable at h, there would be $w \in P_i((h,a'))$ such that $x \rhd_i w$, and—with w possible after a' and x possible after a—no type would find a' to be SOD, contrary to pruning. If $w' \rhd_i x$ then no type would find a to be SOD, which contradicts pruning; we conclude that $w' \not\rhd_i x$. In particular, $x \notin P_i((h,a'))$. If $y \rhd_i w'$, then $y \notin P_i((h,a'))$ because $w' \in \hat{P}_i((h,a'))$ is undominated at (h,a'). Thus, a' would not be SOD for any type, a contradiction to pruning. We conclude that $y \not\rhd_i w'$. If $w' \rhd_i y$, then this and the previously established $w' \not\rhd_i x$, gives us the existence of type $\succ_i$ such that $x \succ_i w' \succ_i y$. This type has no SOD action at h, a contradiction to pruning. We conclude that $w' \not\rhd_i y$. If $x \rhd_i w'$, then type $x \succ_i w' \succ_i y$ exists and has no SOD action at h; we conclude that $x \not\rhd_i w'$. The above four conclusions imply that x, y, w' are mutually undominated at h. Thus, there is a type such that $x \succ_i w' \succ_i y$ and this type has no SOD action at h. This final contradiction shows that $\hat{P}_i((h,a')) = \{y\}$ for all $a' \neq a$.

We further claim 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 $y \rhd_i w'$. As the mechanism is pruned, some type $\succ'_i$ takes action a'; but, the worst case from a', for all types, is at best w', while y is possible following a; thus a' is not SOD for type $\succ'_i$. This contradiction shows that $P_i((h,a')) = \{y\}$ for all $a' \neq a$.

Finally, 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 $a' \neq a$. As $\hat{P}_i(h) = \{x,y\}$, z dominates neither x nor y, and richness gives us a type such that $x \succ_i y \succ_i z$. This type has no SOD action at h; this contradicts the mechanism being SOSP and established the theorem.) ∎