

MPRA

Munich Personal RePEc Archive

Obviously Strategy-Proof Mechanisms

Li, Shengwu

Harvard University

2017

Online at <https://mpra.ub.uni-muenchen.de/78930/>

MPRA Paper No. 78930, posted 04 May 2017 17:24 UTC

Obviously Strategy-Proof Mechanisms*

Shengwu Li[†]

<http://ssrn.com/abstract=2560028>

First uploaded: 3 February 2015. This version: 17 April 2017.

Abstract

A strategy is *obviously dominant* if, for any deviation, at any information set where both strategies first diverge, the best outcome under the deviation is no better than the worst outcome under the dominant strategy. A mechanism is obviously strategy-proof (OSP) if it has an equilibrium in obviously dominant strategies. This has a behavioral interpretation: A strategy is obviously dominant iff a cognitively limited agent can recognize it as weakly dominant. It also has a classical interpretation: A choice rule is OSP-implementable iff it can be carried out by a social planner under a particular regime of partial commitment.

1 Introduction

Dominant-strategy mechanisms are often said to be desirable. They reduce participation costs and cognitive costs, by making it easy for agents to decide what to do.¹ They protect agents from strategic errors.² Dominant-strategy mechanisms prevent waste from rent-seeking espionage, since spying on other players yields no strategic advantage.

*I thank especially my advisors, Paul Milgrom and Muriel Niederle. I thank Nick Arnosti, Roland Benabou, Douglas Bernheim, Gabriel Carroll, Paul J. Healy, Matthew Jackson, Fuhito Kojima, Roger Myerson, Michael Ostrovsky, Matthew Rabin, Alvin Roth, Ilya Segal, and four anonymous referees for their invaluable advice. I thank Paul J. Healy for his generosity in allowing my use of the Ohio State University Experimental Economics Laboratory, and Luyao Zhang for her help in running the experiments. This work was supported by the Kohlhaugen Fellowship Fund, through a grant to the Stanford Institute for Economic Policy Research. All errors remain my own.

[†]shengwu.li@fas.harvard.edu

¹Vickrey (1961) writes that, in second-price auctions: “Each bidder can confine his efforts and attention to an appraisal of the value the article would have in his own hands, at a considerable saving in mental strain and possibly in out-of-pocket expense.”

²For instance, school choice mechanisms that lack dominant strategies may harm parents who do not strategize well (Pathak and Sönmez, 2008).

Moreover, the resulting outcome does not depend sensitively on each agent’s higher-order beliefs.

These benefits largely depend on agents *understanding* that the mechanism has an equilibrium in dominant strategies; *i.e.* that it is strategy-proof (SP). Only then can they conclude that they need not attempt to discover their opponents’ strategies or to game the system.³

However, some strategy-proof mechanisms are simpler for real people to understand than others. For instance, choosing when to quit in an ascending clock auction is the same as choosing a bid in a second-price sealed-bid auction (Vickrey, 1961). The two formats are strategically equivalent; they have the same reduced normal form.⁴ Nonetheless, laboratory subjects are substantially more likely to play the dominant strategy under a clock auction than under sealed bids (Kagel et al., 1987). Theorists have also expressed this intuition:

Some other possible advantages of dynamic auctions over static auctions are difficult to model explicitly within standard economics or game-theory frameworks. For example, . . . it is generally held that the English auction is simpler for real-world bidders to understand than the sealed-bid second-price auction, leading the English auction to perform more closely to theory. (Ausubel, 2004)

In this paper, I model what it means for a mechanism to be *obviously* strategy-proof. This approach invokes no new primitives. Thus, it identifies a set of mechanisms as simple to understand, while remaining as parsimonious as standard game theory.

A strategy S_i is *obviously dominant* if, for any deviating strategy S'_i , starting from any earliest information set where S_i and S'_i diverge, the best possible outcome from S'_i is no better than the worst possible outcome from S_i . A mechanism is *obviously strategy-proof* (OSP) if it has an equilibrium in obviously dominant strategies. By construction, OSP depends on the extensive game form, so two games with the same normal form may differ on this criterion. Obvious dominance implies weak dominance, so OSP implies SP.

This definition distinguishes ascending auctions and second-price sealed-bid auctions. Ascending auctions are obviously strategy-proof. Suppose you value the object at \$10. If the current price is below \$10, then the best possible outcome from quitting now is

³ Policymakers could announce that a mechanism is strategy-proof, but that may not be enough. If agents do not understand the mechanism well, then they may be justifiably skeptical of such declarations. For instance, Google’s advertising materials for the Generalized Second-Price auction appeared to imply that it was strategy-proof, when in fact it was not (Edelman et al., 2007).

⁴ This equivalence assumes that we restrict attention to cut-off strategies in ascending auctions.

no better than the worst possible outcome from staying in the auction (and quitting at \$10). If the price is above \$10, then the best possible outcome from staying in the auction is no better than the worst possible outcome from quitting now.

Second-price sealed-bid auctions are strategy-proof, but not obviously strategy-proof. Consider the strategies “bid \$10” and “bid \$11”. The earliest information set where these diverge is the point where you submit your bid. If you bid \$11, you might win the object at some price strictly below \$10. If you bid \$10, you might not win the object. The best possible outcome from deviating is better than the worst possible outcome from truth-telling. This captures an intuition expressed by experimental economists:

The idea that bidding modestly in excess of x only increases the chance of winning the auction when you don’t want to win is far from obvious from the sealed bid procedure. (Kagel et al., 1987)

I produce two characterization theorems, which suggest two interpretations of obvious strategy-proofness. The first interpretation is behavioral: Obviously dominant strategies are those that can be recognized as dominant by a cognitively limited agent. The second interpretation is classical: Obviously strategy-proof mechanisms are those that can be carried out by a social planner with only partial commitment power.


First, I model an agent who has a simplified mental representation of the world: Instead of understanding every detail of every game, his understanding is limited by a coarse partition on the space of all games. I show that a strategy S_i is obviously dominant if and only if such an agent can recognize S_i as weakly dominant.

Consider the mechanisms in Figure 1. Suppose Agent 1 has preferences: $A \succ B \succ C \succ D$. In mechanism (i), it is a weakly dominant strategy for 1 to play L . Both mechanisms are intuitively similar, but it is not a weakly dominant strategy for Agent 1 to play L in mechanism (ii).

In order for Agent 1 to recognize that it is weakly dominant to play L in mechanism (i), he must use contingent reasoning. That is, he must think through hypothetical scenarios: “If Agent 2 plays l , then I should play L , since I prefer A to B . If Agent 2 plays r , then I should play L , since I prefer C to D . Therefore, I should play L , no matter what Agent 2 plays.” Notice that the quoted inferences are valid in (i), but not valid in (ii).

Suppose Agent 1 is unable to engage in contingent reasoning. That is, he knows that playing L might lead to A or C , and playing R might lead to B or D . However, he does not understand how, state-by-state, the outcomes after playing L are related to the outcomes after playing R . Then it is as though he cannot distinguish (i) and (ii).

Figure 1: ‘Similar’ mechanisms from 1’s perspective.



This idea can be made formal and general. I define an equivalence relation on the space of mechanisms: The *experience* of agent i at history h records the information sets where i was called to play, and the actions that i took, in chronological order.⁵ Two mechanisms G and G' are i -*indistinguishable* if there is a bijection from i 's information sets and actions in G , onto i 's information sets and actions in G' , such that:

1. G can produce for i some experience if and only if G' can produce for i the image (under the bijection) of that experience.
2. An experience might result in some outcome in G if and only if its image might result in that same outcome in G' .

With this relation, we can partition the set of all mechanisms into equivalence classes. For instance, the mechanisms in Figure 1 are 1-indistinguishable.

Our agent knows the experiences that a mechanism might generate, and the resulting outcomes. Thus, at any of his information sets, for any continuation strategy, he knows all the outcomes that might result from that strategy. However, this does not nail down every detail of the mechanism. In particular, our agent is unable to reason contingently about hypothetical scenarios. At any information set, if S_i and S'_i prescribe different actions, our agent is unable to make inferences of the form, “If playing S_i leads to outcome x , then playing S'_i leads to outcome x' .”

For instance, in a second-price auction, our agent knows that for each bid, he will either win the object at a price weakly less than that bid, or not win (and pay zero).

⁵An experience is a standard concept in the theory of extensive games; experiences are sometimes used to define perfect recall.

However, he cannot make the inference, “If bidding \$11 would cause me to win the object at \$8, then bidding \$10 would cause me to win the object at \$8.” Since the agent is unable to make these inferences about the second-price auction, he is unable to correctly identify his dominant strategy.

In general, what strategies can our agent recognize as weakly dominant, when he is constrained by this coarse partition? These are all (and only) the obviously dominant strategies. The first characterization theorem states: A strategy S_i is obviously dominant in G if and only if it is weakly dominant in every G' that is i -indistinguishable from G .

The second characterization theorem for OSP relates to the problem of mechanism design under partial commitment. In mechanism design, we usually assume that the Planner can commit to every detail of a mechanism, including the events that an individual agent does not directly observe. For instance, in a sealed-bid auction, we assume that the Planner can commit to the function from all bid profiles to allocations and payments, even though each agent only directly observes his own bid. Sometimes this assumption is too strong. If agents cannot individually verify the details of a mechanism, the Planner may be unable to commit to it.

Mechanism design under partial commitment is a pressing problem. Auctions run by central brokers over the Internet account for billions of dollars of economic activity (Edelman et al., 2007). In such settings, bidders may be unable to verify that the other bidders exist, let alone what actions they have taken. As another example, some wireless spectrum auctions use computationally demanding techniques to solve complex assignment problems, and the auctioneer may not be permitted to publicly disclose all the bids. In these settings, individual bidders may find it difficult and costly to verify the output of the auctioneer’s algorithm (Milgrom and Segal, 2015).

For the second characterization theorem, I consider a ‘metagame’ where the Planner privately communicates with agents, and eventually decides on an outcome. The Planner chooses one agent, and sends a private message, along with a set of acceptable replies. That agent chooses a reply, which the Planner observes. The Planner can then either repeat this process (possibly with a different agent) or announce an outcome and end the game.

The Planner has partial commitment power: For each agent, she can commit to use only a subset of her available strategies. However, the subset she promises to Agent i must be measurable with respect to i ’s observations in the game. That is, if the Planner plays a strategy not in that subset, then there exists some agent strategy profile such that Agent i detects (with certainty) that the Planner has deviated. We call this a *bilateral commitment*.

Suppose we require that each agent i 's strategy be optimal, for any strategies of the other agents, and for any Planner strategy compatible with the commitment made to i . What choice rules can be implemented in this metagame?

The second characterization theorem states: A choice rule can be supported by bilateral commitments if and only if that choice rule is OSP-implementable. Consequently, in addition to formalizing a notion of cognitive simplicity, OSP also captures the set of choice rules that can be carried out with only bilateral commitments.

After defining and characterizing OSP, I apply this concept to several mechanism design environments.

For the first application, I consider binary allocation problems. In this environment, there is a set of agents N with continuous single-dimensional types $\theta_i \in [\underline{\theta}_i, \bar{\theta}_i]$. An *allocation* y is a subset of N . An *allocation rule* f_y is a function from type profiles to allocations. We augment this with a *transfer rule* f_t , which specifies money transfers for each agent. Each agent has utility equal to his type if he is in the allocation, plus his net transfer.

$$u_i(\theta_i, y, t) = 1_{i \in y} \theta_i + t_i \tag{1}$$

Binary allocation problems encompass several canonical settings. They include private-value auctions with unit demand. They include procurement auctions with unit supply; not being in the allocation is ‘winning the contract’, and the bidder’s type is his cost of provision. They also include binary public good problems; the feasible allocations are N and the empty set.

Mechanism design theory has extensively investigated SP-implementation in this environment. f_y is SP-implementable if and only if f_y is monotone in each agent’s type (Spence, 1974; Mirrlees, 1971; Myerson, 1981). If f_y is SP-implementable, then the required transfer rule f_t is essentially unique (Green and Laffont, 1977; Holmström, 1979).

What are analogues of these canonical results, if we require OSP-implementation rather than SP-implementation? Are ascending clock auctions special, or are there other OSP mechanisms in this environment?

I prove the following theorem: Every mechanism that OSP-implements an allocation rule is ‘essentially’ a *personal-clock auction*, which is a new generalization of ascending auctions. Moreover, this is a full characterization of OSP mechanisms: For any personal-clock auction, there exists some allocation rule that it OSP-implements.

These results imply that when we desire OSP-implementation in a binary allocation problem, we need not search the space of all extensive game forms. Without loss of

generality, we can focus our attention on the class of personal-clock auctions.⁶

As a second application, I produce an impossibility result for a classic matching algorithm: With 3 or more agents, there does not exist a mechanism that OSP-implements Top Trading Cycles (Shapley and Scarf, 1974).

I conduct a laboratory experiment to test the theory. In the experiment, I compare three pairs of mechanisms. In each pair, both mechanisms implement the same choice rule. One mechanism is obviously strategy-proof. The other mechanism is strategy-proof, but not obviously strategy-proof. Standard theory predicts that both mechanisms result in dominant strategy play, and have identical outcomes. Instead, subjects play the dominant strategy at significantly higher rates under the OSP mechanism, compared to the mechanism that is just SP. This effect occurs for all three pairs of mechanisms, and persists even after playing each mechanism five times with feedback.

The rest of the paper proceeds in the usual order. Section 2 reviews the literature. Section 3 provides formal definitions and characterizations. Section 4 covers applications. Section 5 reports the laboratory experiment. Section 6 concludes. Proofs omitted from the main text are in Appendix B.

2 Related Literature

It is widely acknowledged that ascending auctions are simpler for real bidders than second-price sealed-bid auctions (Ausubel, 2004). Laboratory experiments have investigated and corroborated this claim (Kagel et al., 1987; Kagel and Levin, 1993). More generally, laboratory subjects find it difficult to reason state-by-state about hypothetical scenarios (Charness and Levin, 2009; Esponda and Vespa, 2014; Ngangoue and Weizsacker, 2015). This mental process, often called “contingent reasoning”, has received little formal treatment in economic theory.⁷

There is also a strand of literature, including Vickrey’s seminal paper, that observes that sealed-bid auctions raise problems of commitment (Vickrey, 1961; Cramton, 1998). For instance, it may be difficult to prevent shill bidding without third-party verification. Rothkopf et al. (1990) argue that “robustness in the face of cheating and of fear of cheating is important in determining auction form”.

This paper formalizes and unifies these two strands of thought. It shows that mechanisms that do not require contingent reasoning are identical to mechanisms that can be

⁶Of course, if we do not impose the additional structure of a binary allocation problem, then there exist OSP mechanisms that are not personal-clock auctions. This paper contains several examples.

⁷In subsequent work, Esponda and Vespa (2016) investigate axioms governing contingent reasoning in single-agent decision problems.

run under bilateral commitment.

For combinatorial auctions, Vickrey-Clarke-Groves mechanisms can be strategically complex and computationally infeasible. Consequently, there has been substantial interest in designing ‘simple’ mechanisms that perform well, such as deferred-acceptance clock auctions (Milgrom and Segal, 2015) and posted-price mechanisms (Bartal et al., 2003; Feldman et al., 2014; Dütting et al., 2016). These have the (previously unmodeled) advantage of being obviously strategy-proof.

OSP is distinct from O-solvability, a solution concept used in the computer science literature on decentralized learning (Friedman, 2002, 2004). Strategy S_i *overwhelms* S'_i if the worst possible outcome from S_i is strictly better than the best possible outcome from S'_i . O-solvability calls for the iterated deletion of overwhelmed strategies. One difference between the two concepts is that O-solvability is for normal form games, whereas OSP invokes a notion of an ‘earliest point of departure’, which is only defined in the extensive form. O-solvability is too strong for our current purposes, because almost no games studied in mechanism design are O-solvable.⁸

There is a small literature that builds on this paper. In school choice settings, there exist school priorities such that no mechanism OSP-implements the deferred acceptance algorithm for those priorities (Ashlagi and Gonczarowski, 2015). Bade and Gonczarowski (2016) characterize OSP mechanisms for house matching, and in social choice environments with single-peaked preferences. Pycia and Troyan (2016) characterize OSP mechanisms in general settings with no transfers and ‘rich’ preferences, and propose a stronger solution concept (strong obvious strategy-proofness). Zhang and Levin (2017) provide decision-theoretic foundations for obvious dominance, and propose a weaker solution concept (equilibrium in partition dominant strategies). Li (2017) defines obvious *ex post* equilibrium for settings with interdependent values.

3 Definition and Characterization

The planner operates in an *environment* consisting of:

1. A set of agents, $N \equiv \{1, \dots, n\}$.
2. A set of outcomes, X .
3. A set of type profiles, $\Theta_N \equiv \prod_{i \in N} \Theta_i$.
4. A utility function for each agent, $u_i : X \times \Theta_i \rightarrow \mathbb{R}$

⁸For instance, neither ascending clock auctions nor second-price sealed-bid auctions are O-solvable.

A mechanism is an *extensive game form with consequences in X* .⁹ This is an extensive game form where each terminal history z results in some outcome $g(z) \in X$. We restrict attention to game forms with perfect recall¹⁰ and finite depth¹¹. The full definition is familiar to most readers, so we relegate it to Appendix A. \mathcal{G} denotes the set of all such game forms, with representative element G . Useful notation is compiled in Table 1.

Table 1: Notation for Extensive Game Forms

Name	Notation	Representative element
histories	H	h
precedence relation over histories	\prec	
immediate successors of h	$\sigma(h)$	
initial history	h_\emptyset	
terminal histories	Z	z
outcome resulting from z	$g(z)$	
player (agent or chance) called to play at h	$P(h)$	
information sets for agent i	\mathcal{I}_i	I_i
actions available at I_i	$A(I_i)$	
most recent action at h	$\mathcal{A}(h)$	
probability measure for chance moves	δ_c	
realization of chance moves	d_c	

We write $I_i \prec I'_i$ if there exist histories $h \in I_i$ and $h' \in I'_i$ such that $h \prec h'$. We write $I_i \prec h$ if there exists $h' \in I_i$ such that $h' \prec h$. $h \prec I_i$ is defined symmetrically. We use \preceq to denote the corresponding weak order.

A *strategy* $S_i(\cdot)$ chooses an action at every information set for agent i , $S_i(I_i) \in A(I_i)$. A *strategy profile* $S_N = (S_i)_{i \in N}$ specifies a strategy for each agent.

A *type-strategy* $\mathbf{S}_i(\cdot)$ specifies a strategy for every type of agent i , where $\mathbf{S}_i(\theta_i)$ denotes the strategy assigned to type θ_i . A *type-strategy profile* $\mathbf{S}_N = (\mathbf{S}_i)_{i \in N}$ specifies a type-strategy for each agent.

Let $z^G(h, S_N, \delta_c)$ be the lottery over terminal histories that results in game form G when we start from h and play proceeds according to (S_N, δ_c) . $z^G(h, S_N, d_c)$ is the

⁹By adopting this paradigm, we assume that the mechanism describes the *entire* strategic interaction between the agents. For instance, after the mechanism has concluded, agent 1 cannot send money to agent 2, and agent 2 cannot throw a brick through agent 1's window. Such post-game moves often imply that agents do not have dominant strategies, let alone obviously dominant strategies. Savage (1954) notes that "the use of modest little worlds, tailored to particular contexts, is often a simplification, the advantage of which is justified by a considerable body of mathematical experience with related ideas."

¹⁰ G has *perfect recall* if for any information set I_i , for any two histories h and h' in I_i , $\psi_i(h) = \psi_i(h')$, where $\psi_i(h)$ is the experience of agent i at history h (see Definition 8).

¹¹That is, for each game form, there exists some number k such that no history has more than k predecessors. This restriction is not essential to the main results, but simplifies the metagame in Theorem 2.

result of one realization of the chance moves under δ_c . We sometimes write this as $z^G(h, S_i, S_{-i}, d_c)$.

Let $u_i^G(h, S_i, S_{-i}, d_c, \theta_i) \equiv u_i(g(z^G(h, S_i, S_{-i}, d_c)), \theta_i)$. This is the utility to agent i in game form G , when we start at history h , play proceeds according to (S_i, S_{-i}, d_c) , and the resulting outcome is evaluated according to preferences θ_i .

Definition 1. Given G and θ_i , S_i is **weakly dominant** if: $\forall S'_i : \forall S_{-i} :$

$$\mathbb{E}_{\delta_c}[u_i^G(h_\emptyset, S'_i, S_{-i}, d_c, \theta_i)] \leq \mathbb{E}_{\delta_c}[u_i^G(h_\emptyset, S_i, S_{-i}, d_c, \theta_i)] \quad (2)$$

Let $\alpha(S_i, S'_i)$ be the set of earliest points of departure for S_i and S'_i . That is, $\alpha(S_i, S'_i)$ contains the information sets where S_i and S'_i choose different actions, that can be on the path of play given S_i and given S'_i .

Definition 2 (Earliest Points of Departure). $I_i \in \alpha(S_i, S'_i)$ if and only if:

1. $S_i(I_i) \neq S'_i(I_i)$
2. There exist S_{-i}, d_c such that $I_i \prec z^G(h_\emptyset, S_i, S_{-i}, d_c)$.
3. There exist S_{-i}, d_c such that $I_i \prec z^G(h_\emptyset, S'_i, S_{-i}, d_c)$.

This definition can be extended to deal with mixed strategies¹², but pure strategies suffice for our current purposes.

Definition 3. Given G and θ_i , S_i is **obviously dominant** if: $\forall S'_i : \forall I_i \in \alpha(S_i, S'_i) :$

$$\sup_{h \in I_i, S_{-i}, d_c} u_i^G(h, S'_i, S_{-i}, d_c, \theta_i) \leq \inf_{h \in I_i, S_{-i}, d_c} u_i^G(h, S_i, S_{-i}, d_c, \theta_i) \quad (3)$$

In words, S_i is obviously dominant if, for any deviating strategy S'_i , conditional on reaching any earliest point of departure, the best possible outcome under S'_i is no better than the worst possible outcome under S_i .¹³

Compare Definition 1 and Definition 3. Weak dominance is defined using h_\emptyset , the history that begins the game. Consequently, if two extensive games have the same normal form, then they have the same weakly dominant strategies. Obvious dominance

¹²Three modifications are necessary: First, we change requirement 1 to be that both strategies specify different *probability measures* at I_i . Second, we adapt requirements 2 and 3 to hold for *some realization* of the mixed strategies. Finally, we include the recursive requirement, “There does not exist $I'_i \prec I_i$ such that $I'_i \in \alpha(S_i, S'_i)$.”

¹³Obvious dominance is related to conditional dominance (Shimoji and Watson, 1998); S_i conditionally dominates S'_i at I_i if, for any S_{-i} consistent with reaching I_i , the payoff under S'_i is no better than the payoff under S_i .

is defined with histories that are in information sets that are earliest points of departure. Thus two extensive games with the same normal form may not have the same obviously dominant strategies.¹⁴ Switching to a direct revelation mechanism may not preserve obvious dominance, so the standard revelation principle does not apply.¹⁵

Weak dominance treats chance moves and other players asymmetrically. Suppose G has a Bayes-Nash equilibrium, and we replace all the agents in $N \setminus 1$ with chance moves drawn from the Bayes-Nash equilibrium distribution. Then agent 1 has a weakly dominant strategy. By contrast, obvious dominance treats chance moves and other players symmetrically.

A *choice rule* is a function $f : \Theta_N \rightarrow X$. If we consider stochastic choice rules, then it is a function $f : \Theta_N \rightarrow \Delta X$.¹⁶

A *solution concept* $\mathcal{C}(\cdot)$ is a set-valued function; for each G , it specifies a set of type-strategy profiles $\mathcal{C}(G)$, which may be an empty set.

Definition 4. (G, \mathbf{S}_N) *\mathcal{C} -implements* f if:

1. $\mathbf{S}_N \in \mathcal{C}(G)$.
2. $\forall \theta_N \in \Theta_N : f(\theta_N) = g(z^G(h_\emptyset, (\mathbf{S}_i(\theta_i))_{i \in N}, \delta_c))$

G *\mathcal{C} -implements* f if there exists \mathbf{S}_N that satisfies the above requirements.

f is *\mathcal{C} -implementable* if there exist (G, \mathbf{S}_N) that satisfy the above requirements.

Note that our concern is with weak implementation: We require that $\mathbf{S}_N \in \mathcal{C}(G)$, not $\{\mathbf{S}_N\} = \mathcal{C}(G)$. This is to preserve the analogy with canonical results for strategy-proofness, many of which assume weak implementation (Myerson, 1981; Saks and Yu, 2005).

Definition 5 (Strategy-Proof). $\mathbf{S}_N \in \mathbf{SP}(G)$ if for all i and for all θ_i , $\mathbf{S}_i(\theta_i)$ is weakly dominant.

¹⁴Two extensive games have the same reduced normal form if and only if they can be made identical using a small set of elementary transformations (Thompson, 1952; Elmes and Reny, 1994). Which of these transformations does not preserve obvious dominance? Elmes and Reny (1994) propose three such transformations, INT, COA, and ADD, which preserve perfect recall. Brief inspection reveals that obvious dominance is invariant under INT and COA, but varies under ADD.

¹⁵Glazer and Rubinstein (1996) argue that extensive games can be easier to dominance-solve than normal-form games, because backward induction provides guidance about the correct order to delete strategies. The standard revelation principle does not apply to their solution concept either.

¹⁶For readability, we generally suppress the latter notation, but the claims that follow hold for both deterministic and stochastic choice rules. Additionally, the set X could *itself* be a set of lotteries. The interpretation of this is that the planner can carry out one-time public lotteries at the end of the mechanism, where the randomization is observable and verifiable.

Definition 6 (Obviously Strategy-Proof). $\mathbf{S}_N \in \mathbf{OSP}(G)$ if for all i and for all θ_i , $\mathbf{S}_i(\theta_i)$ is obviously dominant.

A mechanism is weakly group-strategy-proof if there does not exist a coalition that could deviate and all be *strictly* better off *ex post*.

Definition 7 (Weakly Group-Strategy-Proof). $\mathbf{S}_N \in \mathbf{WGSP}(G)$ if there does not exist a coalition $\hat{N} \subseteq N$, type profile θ_N , deviating strategies $\hat{S}_{\hat{N}}$, non-coalition strategies $S_{N \setminus \hat{N}}$ and chance moves d_c such that: For all $i \in \hat{N}$:

$$u_i^G(h_\emptyset, \hat{S}_{\hat{N}}, S_{N \setminus \hat{N}}, d_c, \theta_i) > u_i^G(h_\emptyset, \mathbf{S}_{\hat{N}}(\theta_{\hat{N}}), S_{N \setminus \hat{N}}, d_c, \theta_i) \quad (4)$$

Obvious strategy-proofness implies weak group-strategy-proofness.¹⁷

Proposition 1. If $\mathbf{S}_N \in \mathbf{OSP}(G)$, then $\mathbf{S}_N \in \mathbf{WGSP}(G)$.

Proof. Suppose $\mathbf{S}_N \notin \mathbf{WGSP}(G)$. Then there is a coalition \hat{N} with types $\hat{\theta}_{\hat{N}}$ that could jointly deviate to strategies $\hat{S}_{\hat{N}}$ and all be strictly better off. Fix $S_{N \setminus \hat{N}}$ and d_c such that all agents in the coalition are strictly better off. Along the resulting terminal history, there is a first agent i in the coalition to deviate from $\mathbf{S}_i(\theta_i)$ to \hat{S}_i . That first deviation happens at some information set $I_i \in \alpha(\mathbf{S}_i(\theta_i), \hat{S}_i)$. Agent i strictly gains from that deviation, so $\mathbf{S}_N \notin \mathbf{OSP}(G)$.¹⁸ \square

Corollary 1. If $\mathbf{S}_N \in \mathbf{OSP}(G)$, then $\mathbf{S}_N \in \mathbf{SP}(G)$.

Proposition 1 suggests a question: Is a choice rule OSP-implementable if and only if it is WGSP-implementable? Proposition 5 shows that this is not so.

3.1 Cognitive limitations

In what sense is obvious dominance obvious? Intuitively, to see that S'_i is weakly dominated by S_i , the agent must understand the entire function u_i^G , and check that for all opponent strategy profiles S_{-i} , the payoff from S'_i is no better than the payoff from S_i . By contrast, to see that S'_i is obviously dominated by S_i , the agent need only know the *range* of the functions $u_i^G(\cdot, S'_i, \dots)$ and $u_i^G(\cdot, S_i, \dots)$ at any earliest point of departure. Thus, obvious dominance can be recognized even if the agent has a simplified mental model of the world. We now make this point rigorously.

¹⁷Barberà et al. (2016) note that “many well-known individually strategy-proof mechanisms are also group strategy-proof, even if the latter is in principle a much stronger condition than the former.”

¹⁸I thank Ilya Segal for suggesting this concise proof.

Definition 8. $\psi_i(h)$ denotes i 's **experience** along history h ; it is the sequence of information sets where i was called to play, and the actions that i took, in order. We construct this by starting at h_\emptyset , and moving step-by-step through predecessors of h . At each $h' \preceq h$, if $P(h') = i$, then we add the current information set to the sequence. If i has just played an action at h' , then we add that action to the sequence.

We use Ψ_i to denote the set $\{\psi_i(h) \mid h \in H\} \cup \psi_\emptyset$, where ψ_\emptyset is the empty sequence.¹⁹ We define $\psi_i(I_i) := \psi_i(h) \mid h \in I_i$.

We define an equivalence relation between mechanisms. In words, G and G' are i -indistinguishable if there exists a bijection from i 's information sets and actions in G onto i 's information sets and actions in G' , such that:

1. ψ_i is an experience in G iff ψ_i 's image is an experience in G' .
2. Outcome x could follow experience ψ_i in G iff x could follow ψ_i 's image in G' .

Definition 9. Take any $G, G' \in \mathcal{G}$, with information partitions $\mathcal{I}_i, \mathcal{I}'_i$ and experience sets Ψ_i, Ψ'_i . G and G' are i -**indistinguishable** if there exists a bijection $\lambda_{G,G'}$ from $\mathcal{I}_i \cup A(\mathcal{I}_i)$ to $\mathcal{I}'_i \cup A'(\mathcal{I}'_i)$ such that:

1. $\psi_i \in \Psi_i$ iff $\lambda_{G,G'}(\psi_i) \in \Psi'_i$
2. $\exists z \in Z : g(z) = x, \psi_i(z) = \psi_i$ iff $\exists z' \in Z' : g'(z') = x, \psi'_i(z') = \lambda_{G,G'}(\psi_i)$

where we use $\lambda_{G,G'}(\psi_i)$ to denote the sequence produced by passing every element of ψ_i through $\lambda_{G,G'}$.

For G and G' that are i -indistinguishable, we define $\lambda_{G,G'}(S_i)$ to be the strategy that, given information set I'_i in G' , plays $\lambda_{G,G'}(S_i(\lambda_{G,G'}^{-1}(I'_i)))$.

The next theorem states that obviously dominant strategies are the strategies that can be recognized as weakly dominant, by an agent who has this simplified mental model of the world.

Theorem 1. For any i, θ_i : S_i is obviously dominant in G if and only if for every G' that is i -indistinguishable from G , $\lambda_{G,G'}(S_i)$ is weakly dominant in G' .

The “if” direction permits a constructive proof. Suppose S_i is not obviously dominant in G . Then there is some deviating strategy S'_i and some earliest point of departure I_i where the obvious dominance inequality does not hold. We construct a game form

¹⁹Mandating the inclusion of the empty sequence has the following consequence: By looking at the set Ψ_i , it is not possible to infer whether $P(h_\emptyset) = i$.

G' that is i -indistinguishable from G , such that, for some opponent strategies, if i plays $\lambda_{G,G'}(S_i)$ then play proceeds according to the worst-case scenario (consistent with reaching I_i), and if i deviates at $\lambda_{G,G'}(I_i)$, then play proceeds according to some best-case scenario (consistent with reaching I_i). The key is to find a general construction that always remains in the same equivalence class. The “only if” direction proceeds as follows: Suppose there exists some G' in the equivalence class of G , where $\lambda_{G,G'}(S_i)$ is not weakly dominant. There exists some earliest information set in G' where i could gain by deviating. We then use $\lambda_{G,G'}^{-1}$ to locate an information set in G , and a deviation S'_i , that do not satisfy the obvious dominance inequality. Appendix B provides the details.

One interpretation of Theorem 1 is that obviously dominant strategies are those that can be recognized as dominant given only a *partial description* of the game form.

Another interpretation of Theorem 1 is that obviously dominant strategies are robust to *local misunderstandings*. Suppose an agent could mistake any G for any i -indistinguishable G' . He has some belief about how his opponents are playing in G' , and best responds to that (mistaken) belief. For instance, when faced with a second-price auction, he could believe that he faces a *third-price* auction²⁰, in which case bidding above his value is a best response to the (symmetric) Nash equilibrium opponent strategies (Kagel and Levin, 1993). When is his (true) dominant strategy still in his set of best responses, for any such local mistake? By Theorem 1, this holds if and only if it is obviously dominant.

Many solution concepts in behavioral game theory specify that agents understand the game form, but best-respond to mistaken beliefs about opponent strategies.²¹ Since a dominant strategy is always a best response, such theories predict no mistakes in strategy-proof mechanisms.²² Obvious dominance captures misunderstandings about the game form, and is thus orthogonal to these theories.

3.2 Supported by bilateral commitments

Suppose the following *augmented* game form \tilde{G} with consequences in X : As before we have a set of agents N , outcomes X , and preference profiles $\prod_{i \in N} \Theta_i$. However, there is one player in addition to N : Player 0, the Planner.

²⁰In both second-price and third-price auctions, the bidder submits a sealed bid. For any sealed bid, either he receives the object and pays a price less than or equal to his bid, or does not receive the object and makes no payments. Thus, both mechanisms are i -indistinguishable.

²¹Such concepts include level- k equilibrium (Stahl and Wilson, 1994, 1995; Nagel, 1995), cognitive hierarchy equilibrium (Camerer et al., 2004), cursed equilibrium (Eyster and Rabin, 2005; Esponda, 2008), and analogy-based expectation equilibrium (Jehiel, 2005).

²²With the exception of level-0 agents in level- k and cognitive hierarchy models.

The Planner has an arbitrarily rich message space M . At the start of the game, each agent $i \in N$ privately observes θ_i . Play proceeds as follows:

1. The Planner chooses one agent $i \in N$ and sends a query $m \in M$, along with a set of acceptable replies $R \subset M$.²³
2. i observes (m, R) , and chooses a reply $r \in R$.
3. The Planner observes r .
4. The Planner either selects an outcome $x \in X$, or chooses to send another query.
 - (a) If the Planner selects an outcome, the game ends.
 - (b) If the Planner chooses to send another query, go to Step 1.

For $i \in N$, i 's type-strategy specifies what reply to give, as a function of his preferences, the past sequence of queries and replies between him and the Planner, and the current (m, R) . That is:

$$\tilde{\mathbf{S}}_i(\theta_i, (m_k, R_k, r_k)_{k=1}^{t-1}, m_t, R_t) \in R_t \quad (5)$$

Analogously, a strategy \tilde{S}_i depends on $((m_k, R_k, r_k)_{k=1}^{t-1}, m_t, R_t)$, but not on θ_i . We use $\tilde{\mathbf{S}}_i^{\theta_i}$ to denote the strategy played by type θ_i of agent i . We abbreviate $(\tilde{\mathbf{S}}_i^{\theta_i})_{i \in N} \equiv \tilde{\mathbf{S}}_N^{\theta_N}$.

$\tilde{\mathbb{S}}_0$ denotes a pure strategy for the Planner; where \mathbb{S}_0 denotes the set of all pure strategies. We restrict the Planner to send only finitely many queries.

The standard full commitment paradigm is equivalent to allowing the Planner to commit to a unique $\tilde{S}_0 \in \mathbb{S}_0$ (or some probability measure over \mathbb{S}_0). Instead, we assume that for each agent, the Planner can commit to a subset $\hat{\mathbb{S}}_0^i \subseteq \mathbb{S}_0$ that is measurable with respect to that agent's observations in the game.

This is formalized as follows: Each $(\tilde{S}_0, \tilde{S}_N)$ results in some observation $o_i \equiv (o_i^C, o_i^X)$, consisting of a communication sequence between the Planner and agent i , $o_i^C = (m_k, R_k, r_k)_{k=1}^T$ for $T \in \mathbb{N}$, as well as some outcome $o_i^X \in X$.²⁴ \mathcal{O}_i is the set of all possible observations (for agent i). We define $\phi_i : \mathbb{S}_0 \times \mathbb{S}_N \rightarrow \mathcal{O}_i$, where $\phi_i(\tilde{S}_0, \tilde{S}_N)$ is the unique observation resulting from $(\tilde{S}_0, \tilde{S}_N)$. Next we define, for any $\hat{\mathbb{S}}_0 \subseteq \mathbb{S}_0$:

$$\Phi_i(\hat{\mathbb{S}}_0) \equiv \{o_i \mid \exists \tilde{S}_0 \in \hat{\mathbb{S}}_0 : \exists \tilde{S}_N : o_i = \phi_i(\tilde{S}_0, \tilde{S}_N)\} \quad (6)$$

²³This game could be made simpler without altering Theorem 2; the Planner could send only a set of acceptable replies $R \subset M$. Any information contained in the query could simply be 'written into' every acceptable reply. However, distinguishing queries and replies makes the exposition more intuitive.

²⁴The communication sequence might be empty, which we represent using $T = 0$.

For any $\hat{\mathcal{O}}_i \subseteq \mathcal{O}_i$:

$$\Phi_i^{-1}(\hat{\mathcal{O}}_i) \equiv \{\tilde{S}_0 \mid \forall \tilde{S}_N : \phi_i(\tilde{S}_0, \tilde{S}_N) \in \hat{\mathcal{O}}_i\} \quad (7)$$

Definition 10. $\hat{\mathbb{S}}_0$ is *i -measurable* if there exists $\hat{\mathcal{O}}_i$ such that:

$$\hat{\mathbb{S}}_0 = \Phi_i^{-1}(\hat{\mathcal{O}}_i) \quad (8)$$

Intuitively, the i -measurable subsets of \mathbb{S}_0 are those such that, if the Planner deviates, then there exists an agent strategy profile such that agent i detects the deviation; that is, i has an observation that is not compatible with *any* strategy in the promised set. Formally, the i -measurable subsets of \mathbb{S}_0 are the σ -algebra generated by Φ_i (where we impose the discrete σ -algebra on \mathcal{O}_i).²⁵

Definition 11. A *mixed strategy of finite length* over $\hat{\mathbb{S}}_0$ specifies a probability measure over a subset $\tilde{\mathbb{S}}_0 \subseteq \hat{\mathbb{S}}_0$ such that: There exists $k \in \mathbb{N}$ such that: For all $\tilde{S}_0 \in \tilde{\mathbb{S}}_0$ and all $\tilde{S}_N : (\tilde{S}_0, \tilde{S}_N)$ results in the Planner sending k or fewer total queries.

We use $\Delta\hat{\mathbb{S}}_0$ to denote the mixed strategies of finite length over $\hat{\mathbb{S}}_0$. \tilde{S}_0^Δ denotes an element of such a set.

Definition 12. A choice rule f is *supported by bilateral commitments* $(\hat{\mathbb{S}}_0^i)_{i \in N}$ if

1. For all $i \in N$: $\hat{\mathbb{S}}_0^i$ is i -measurable.

2. There exist \tilde{S}_0^Δ , and $\tilde{\mathbf{S}}_N$ such that:

(a) For all $\theta_N : (\tilde{S}_0^\Delta, \tilde{\mathbf{S}}_N^{\theta_N})$ results in $f(\theta_N)$.

(b) $\tilde{S}_0^\Delta \in \Delta \cap_{i \in N} \hat{\mathbb{S}}_0^i$

(c) For all $i \in N$, θ_i , $\tilde{S}'_{N \setminus i}, \tilde{S}_0^{\Delta'} \in \Delta\hat{\mathbb{S}}_0^i$: $\tilde{\mathbf{S}}_i^{\theta_i}$ is a best response to $(\tilde{S}_0^{\Delta'}, \tilde{S}'_{N \setminus i})$ given preferences θ_i .

Requirement 2.a is that the Planner's mixed strategy and the agent's pure strategies result in the (distribution over) outcomes required by the choice rule. Requirement 2.b is that the Planner's strategy is a (possibly degenerate) mixture over pure strategies

²⁵Notice that an observation for player i is simply a sequence of messages and responses, followed by an outcome. It does not include additional information about calendar time. A consequence of this formulation is that the Planner is unable to commit to the order in which she approaches the players.

compatible with *every* bilateral commitment.²⁶ 2.c is that each agent i 's assigned strategy is weakly dominant, when we consider the Planner as a player restricted to playing mixtures over strategies in $\hat{\mathcal{S}}_0^i$.

“Supported by bilateral commitments” is just one of many partial commitment regimes. This one requires that the commitment offered to each agent is measurable with respect to events that he can observe. In reality, contracts are seldom enforceable unless each party can observe breaches. Thus, “supported by bilateral commitments” is a natural case to study.

Theorem 2. *f is OSP-implementable if and only if there exist bilateral commitments $(\hat{\mathcal{S}}_0^i)_{i \in N}$ that support f .*

The intuition behind the proof is as follows: A bilateral commitment $\hat{\mathcal{S}}_0^i$ is essentially equivalent to the Planner committing to ‘run’ only games in some i -indistinguishable equivalence class of \mathcal{G} . Consequently, we can find a set of bilateral commitments that support f_y if and only if we can find some (G, \mathbf{S}_N) such that, for every i , for every θ_i , for every G' that is i -indistinguishable from G , $\lambda_{G, G'}(\mathbf{S}_i(\theta_i))$ is weakly dominant in G' . By Theorem 1, this holds if and only if f_y is OSP-implementable. Appendix B provides the details.

3.3 The pruning principle

When seeking SP-implementation, we can without loss of generality restrict attention to the class of direct revelation mechanisms, by the revelation principle. The standard revelation principle does not hold for OSP mechanisms. As the second-price auction illustrates, there exist OSP-implementable choice rules that cannot be implemented via a direct revelation mechanism.

However, there is a weaker principle that substantially simplifies the analysis. Here we define the *pruning* of a mechanism with respect to a type-strategy profile. If, for all type profiles, a history is never reached, then we delete that history (and redefine the other parts of the tuple so that what remains is well-formed).

Definition 13 (Pruning). *Take any $G = \langle H, \prec, A, \mathcal{A}, P, \delta_c, (\mathcal{I}_i)_{i \in N}, g \rangle$, and \mathbf{S}_N . $\mathcal{P}(G, \mathbf{S}_N) \equiv \langle \tilde{H}, \tilde{\prec}, \tilde{A}, \tilde{\mathcal{A}}, \tilde{P}, \tilde{\delta}_c, (\tilde{\mathcal{I}}_i)_{i \in N}, \tilde{g} \rangle$ is the **pruning** of G with respect to \mathbf{S}_N , constructed as follows:*

1. $\tilde{H} = \{h \in H \mid \exists \theta_N : \exists d_c : h \preceq z^G(h_\emptyset, (\mathbf{S}_i(\theta_i))_{i \in N}, d_c)\}$

²⁶This requirement prevents the Planner from extracting arbitrary information by making promises and breaking them. Otherwise, the Planner could promise to implement some constant outcome, ask every agent to report their type, and then do whatever she pleased with that information.

2. For all i , if $I_i \in \mathcal{I}_i$ then $(I_i \cap \tilde{H}) \in \tilde{\mathcal{I}}_i$.²⁷
3. $(\tilde{\prec}, \tilde{A}, \tilde{\mathcal{A}}, \tilde{P}, \tilde{\delta}_c, \tilde{g})$ are $(\prec, A, \mathcal{A}, P, \delta_c, g)$ restricted to \tilde{H} .

It turns out that, if some mechanism OSP-implements a choice rule, then the pruning of that mechanism with respect to the equilibrium strategies OSP-implements that same choice rule. Thus, while we cannot restrict attention to direct revelation mechanisms, we can restrict attention to ‘minimal’ mechanisms, such that no histories are off the path of play. This is used both in this paper and in the subsequent literature²⁸ to state clean results.

Proposition 2 (The Pruning Principle). *Let $\tilde{G} \equiv \mathcal{P}(G, \mathbf{S}_N)$, and $\tilde{\mathbf{S}}_N$ be \mathbf{S}_N restricted to \tilde{G} . If (G, \mathbf{S}_N) OSP-implements f , then $(\tilde{G}, \tilde{\mathbf{S}}_N)$ OSP-implements f .*

4 Applications

4.1 Binary Allocation Problems

We now consider a canonical environment, $(N, X, \Theta_N, (u_i)_{i \in N})$. Let $Y \subseteq 2^N$ be the set of feasible *allocations*, with representative element $y \in Y$. An outcome consists of an allocation $y \in Y$ and a transfer for each agent, $X = Y \times \mathbb{R}^n$. $t \equiv (t_i)_{i \in N}$ denotes a profile of transfers.

Preferences are quasilinear. $\Theta_N = \prod_{i \in N} \Theta_i$, where $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i]$, for $0 \leq \underline{\theta}_i < \bar{\theta}_i < \infty$. For $\theta_i \in \Theta_i$,

$$u_i(\theta_i, y, t) = 1_{i \in y} \theta_i + t_i \quad (9)$$

For instance, in a private value auction with unit demand, $i \in y$ if and only if agent i receives at least one unit of the good under allocation y . In a procurement auction, $i \in y$ if and only if i does not incur costs of provision under allocation y . θ_i is agent i 's cost of provision (equivalently, benefit of non-provision). In a binary public goods game, $Y = \{\emptyset, N\}$.

An *allocation rule* is a function $f_y : \Theta \rightarrow Y$. A choice rule is thus a combination of an allocation rule and a payment rule, $f = (f_y, f_t)$, where $f_t : \Theta \rightarrow \mathbb{R}^n$. Similarly, for each game form G , we disaggregate the outcome function, $g = (g_y, g_t)$. In this part, we concern ourselves only with deterministic allocation rules and payment rules, and thus suppress notation involving δ_c and d_c .

²⁷Note that the initial history h_\emptyset is distinct from the empty set. That is to say, $(I_i \cap \tilde{H}) = \emptyset$ does *not* entail that $\{h_\emptyset\} \in \tilde{\mathcal{I}}_i$.

²⁸Ashlagi and Gonczarowski (2015); Bade and Gonczarowski (2016); Pycia and Troyan (2016).

Definition 14. An allocation rule f_y is \mathcal{C} -implementable if there exists f_t such that (f_y, f_t) is \mathcal{C} -implementable. G \mathcal{C} -implements f_y if there exists f_t such that G \mathcal{C} -implements (f_y, f_t)

In a binary allocation problem, f_y is SP-implementable if and only if f_y is monotone.²⁹ This result is implicit in Spence (1974) and Mirrlees (1971), and is proved explicitly in Myerson (1981).³⁰ Moreover, if an allocation rule f_y is SP-implementable, then the accompanying transfer rule f_t is essentially unique.

$$f_{t,i}(\theta_i, \theta_{-i}) = -1_{i \in f_y(\theta)} \inf\{\theta'_i \mid i \in f_y(\theta'_i, \theta_{-i})\} + r_i(\theta_{-i}) \quad (10)$$

where r_i is some arbitrary function of the other agents' preferences. This follows easily by arguments similar to those in Green and Laffont (1977) and Holmström (1979).

In this standard environment, what happens when we require OSP-implementation? In particular, what restrictions does OSP-implementation place on the transfer rules and the extensive game form?

In binary allocation problems, every OSP mechanism is essentially a *personal-clock auction*. It is personal in two respects: Firstly, the clock price and closing rule could vary across agents. Secondly, the clock price could be associated with different consequences for different agents. The full statement of the definition is novel, but we will build it out of familiar parts.

Consider how the ascending auction appears to a single bidder: At each point, there is a 'going transfer' associated with being in the allocation (winning the object). The agent either plays *quit* or *continue*; if he plays *quit*, then he is out of the allocation and makes no payments. If he plays *continue*, then either he is in the allocation and has the going transfer, or the going transfer falls (the going price rises) and he faces the same decision again.

We can generalize this procedure a bit, while still ensuring that the agent has an obviously dominant strategy. The going transfer could fall in any increments, and how much it falls could depend on other agents' actions. The transfer associated with being out of the allocation could be non-zero, as long as it is some (known) fixed amount. Even if the agent chooses *continue*, we could sometimes mandate that he quits, in which case he is out of the allocation (and receives the fixed transfer). The agent need not choose

²⁹ f_y is monotone if for all i , for all θ_{-i} , $1_{i \in f_y(\theta_i, \theta_{-i})}$ is weakly increasing in θ_i .

³⁰ These monotonicity results are for weak SP-implementation rather than full SP-implementation. Weak SP-implementation requires $\mathbf{S}_N \in \text{SP}(G)$. Full SP-implementation requires $\mathbf{S}_N = \text{SP}(G)$. There are monotone allocation rules for which the latter requirement cannot be satisfied. For example, suppose two agents with unit demand. Agent 1 receives one unit iff $v_1 > .5$. Agent 2 receives one unit iff $v_2 > v_1$.

between *continue* and *quit* at every information set; he need only be offered a choice when the going transfer strictly falls. There could be multiple actions that *quit*, all having the same consequence for him.³¹ There could be multiple actions that *continue*, provided that the going transfer will not fall in future and at least one such action guarantees that he is in the allocation. The agent could receive arbitrary information about the history of play. We call this generalized procedure *In-Transfer Falls*.

Consider how a descending-price procurement auction appears to a single supplier (holding one unit of an indivisible good): At each point, there is a ‘going transfer’ associated with being *out* of the allocation. The agent either plays *quit* or *continue*; if he plays *quit*, then he is in the allocation (keeps the good) and receives nothing. If he plays *continue*, then either he is out of the allocation (sells the good) and receives the going transfer, or the going transfer falls and he faces the same decision again. Notice that now the transfer associated with being in the allocation is fixed and the transfer associated with being out of the allocation falls monotonically. We could generalize this in the same ways as before, and call the resulting procedure *Out-Transfer Falls*.

In a *personal-clock auction*, starting from any point an agent first has a non-singleton information set, that agent either faces *In-Transfer Falls* or *Out-Transfer Falls*.

Definition 15. *G* is a **personal-clock auction** if, for every $i \in N$, at every earliest information set I_i^* such that $|A(I_i^*)| > 1$:

1. **(In-Transfer Falls):** There exists a fixed transfer $\bar{t}_i \in \mathbb{R}$, a going transfer $\tilde{t}_i : \{I_i \mid I_i^* \preceq I_i\} \rightarrow \mathbb{R}$, and a set of ‘quitting’ actions A^q such that:

(a) For all z where $I_i^* \prec z$:

i. Either: $i \notin g_y(z)$ and $g_{t,i}(z) = \bar{t}_i$.

ii. Or: $i \in g_y(z)$ and

$$g_{t,i}(z) = \inf_{I_i: I_i^* \preceq I_i \prec z} \tilde{t}_i(I_i) \quad (11)$$

(b) For all $a \in A^q$, for all z such that $a \in \psi_i(z)$: $i \notin g_y(z)$

(c) $A^q \cap A(I_i^*) \neq \emptyset$.

(d) For all $I_i', I_i'' \in \{I_i \mid I_i^* \preceq I_i\}$:

i. If $I_i' \prec I_i''$, then $\tilde{t}_i(I_i') \geq \tilde{t}_i(I_i'')$.

ii. If $I_i' \prec I_i''$, $\tilde{t}_i(I_i') > \tilde{t}_i(I_i'')$, and there does not exist I_i''' such that $I_i' \prec I_i''' \prec I_i''$, then $A^q \cap A(I_i'') \neq \emptyset$.

³¹Although they could affect what happens to other agents - it is incentive compatible for the agent to reveal any information about his type at the point when he quits, and the allocation rule could in principle depend on this information.

- iii. If $I'_i \prec I''_i$ and $\tilde{t}_i(I'_i) > \tilde{t}_i(I''_i)$, then $|A(I'_i) \setminus A^q| = 1$.
- iv. If $|A(I'_i) \setminus A^q| > 1$, then there exists $a \in A(I'_i)$ such that: For all z such that $a \in \psi_i(z): i \in g_y(z)$.

2. Or (**Out-Transfer Falls**): As above, but we substitute every instance of “ $i \in g_y(z)$ ” with “ $i \notin g_y(z)$ ” and vice versa.

Notice what this definition does not require. The going transfer need not be equal across agents. Whether and how much one agent’s going transfer changes could depend on other agents’ actions. Some agents could face In-Transfer Falls, and other agents could face Out-Transfer Falls (a two-sided clock auction).³² Which procedure an agent faces could even depend on other agents’ past actions.

Theorem 3. *If (G, \mathbf{S}_N) OSP-implements f_y , then $\mathcal{P}(G, \mathbf{S}_N)$ is a personal-clock auction. If G is a personal-clock auction, then there exist \mathbf{S}_N and f_y such that (G, \mathbf{S}_N) OSP-implements f_y .*

For any normal-form mechanism, there are typically many equivalent extensive forms. The theory of mechanism design seldom provides general criteria to choose between them. In binary allocation problems, obvious dominance pins down many extensive-form details of the ascending auction, and provides an answer to the question, “Why are ascending auctions so common?”

4.2 Top Trading Cycles

We now produce an impossibility result in a classic matching environment (Shapley and Scarf, 1974). There are n agents in the market, each endowed with an indivisible good. An agent’s type is a vector $\theta_i \in \mathbb{R}^n$. Θ_N is the set of all n by n matrices of real numbers. An outcome assigns one object to each agent. If agent i is assigned object k , he has utility θ_i^k . There are no money transfers.

Given preferences θ and agents $R \subseteq N$, a *top trading cycle* is a set $\emptyset \subset R' \subseteq R$ whose members can be indexed in a cyclic order:

$$R' = \{i_1, i_2, \dots, i_r = i_0\} \tag{12}$$

such that each agent i_k likes i_{k+1} ’s good at least as much as any other good in R . Following Roth (1982), we assume that the algorithm in question has an arbitrary, fixed way of resolving ties.

³²Loertscher and Marx (2017) investigate two-sided clock auctions that are prior-free, asymptotically optimal, and obviously strategy-proof.

Definition 16. f is a top trading cycle rule if, for all θ , $f(\theta)$ is equal to the output of the following algorithm:

1. Set $R^1 := N$
2. For $l = 1, 2, \dots$:
 - (a) Choose some top trading cycle $R' \subseteq R^l$.
 - (b) Carry out the indicated trades.
 - (c) Set $R^{l+1} := R^l \setminus R'$.
 - (d) Terminate if $R^{l+1} = \emptyset$.

The above algorithm is of economic interest, because it finds a core allocation in an economy with indivisible goods (Shapley and Scarf, 1974).

Proposition 3. *If f is a top trading cycle rule, then there exists G that SP-implements f . (Roth, 1982)*

Proposition 4. *If f is a top trading cycle rule, then there exists G that WGSP-implements f . (Bird, 1984)*

Proposition 5. *If f is a top trading cycle rule and $n \geq 3$, then there does not exist G that OSP-implements f .*

Proof. OSP-implementability is a hereditary property of functions. That is, if f is OSP-implementable given domain Θ_N , then the subfunction $f' = f$ with domain $\Theta'_N \subseteq \Theta_N$ is OSP-implementable. Thus, to prove Proposition 5, it suffices to produce a subfunction that is not OSP-implementable.

Consider the following subset $\Theta'_N \subset \Theta_N$. Take agents a, b, c , with endowed goods A, B, C . a has only two possible types, θ_a and θ'_a , such that

$$\begin{aligned} \text{Either } B \succ_a C \succ_a A \succ_a \dots \\ \text{or } C \succ_a B \succ_a A \succ_a \dots \end{aligned} \tag{13}$$

We make the symmetric assumption for b and c .

We now argue by contradiction. Take any G pruned with respect to the truthful strategy profiles, such that (by Proposition 2) G OSP-implements $f' = f$ for domain Θ' . Consider some history h at which $P(h) = a$ with a non-singleton action set. This cannot come before all such histories for b and c .

Suppose not, and suppose $B \succ_a C$. If a chooses the action corresponding to $B \succ_a C$, and faces opponent strategies corresponding to $C \succ_b A$ and $B \succ_c A$, then a receives good A . If a chooses the action corresponding to $C \succ_a B$, and faces opponent strategies corresponding $A \succ_c B$, then a receives good C . Thus, it is not an obviously dominant strategy to choose the action corresponding to $B \succ_a C$. So a cannot be the first to have a non-singleton action set.

By symmetry, this argument applies to b and c as well. So all of the action sets for a , b , and c are singletons, and G does not OSP-implement f' , a contradiction. \square

Proposition 5 implies that the OSP-implementable choice rules are not identical to the WGSP-implementable choice rules.

5 Laboratory Experiment

Are obviously strategy-proof mechanisms easier for real people to understand? The following laboratory experiment provides a straightforward test: We compare pairs of mechanisms that implement the same choice rule. One mechanism in each pair is SP, but not OSP. The other mechanism is OSP. Standard game theory predicts that both mechanisms will produce the same outcome. We are interested in whether subjects play the dominant strategy at higher rates under OSP mechanisms.

5.1 Experiment Design

The experiment is an across-subjects design, comparing three pairs of games. There are four players in each game.

For the first pair, we compare the second-price auction (**2P**) and the ascending clock auction (**AC**). In both these games, subjects bid for a money prize. Subjects have induced affiliated private values; if a subject wins the prize, he earns an amount equal to the value of the prize, minus his payments from the auction. For each subject, his value for the prize is equal to a group draw plus a private adjustment. The group draw is uniformly distributed between \$10 and \$110. The private adjustment is uniformly distributed between \$0 and \$20. All money amounts in these games are in 25-cent increments. Each subject knows his own value, but not the group draw or the private adjustment.³³

³³I use affiliated private values for two reasons. First, in strategy-proof auctions with independent private values, incentives for truthful bidding are weak for bidders with values near the extremes. Affiliation strengthens incentives for these bidders. Second, Kagel et al. (1987) use affiliated private values, and the first part of the experiment is designed to replicate their results.

2P is SP, but not OSP. In **2P**, subjects submit their bids simultaneously. The highest bidder wins the prize, and makes a payment equal to the second-highest bid. Bids are constrained to be between \$0 and \$150.³⁴

AC is OSP. In **AC**, the price starts at a low value (the highest \$25 increment that is below the group draw), and counts upwards, up to a maximum of \$150. Each bidder can quit at any point.³⁵ When only one bidder is left, that bidder wins the object at the current price.

Previous studies comparing second-price auctions to ascending clock auctions have small sample sizes, given that when the same subjects play a sequence of auctions, these are plainly not independent observations. Kagel et al. (1987) compare 2 groups playing second-price auctions to 2 groups playing ascending clock auctions. Harstad (2000) compares 5 groups playing second-price auctions to 3 groups playing ascending clock auctions. (The comparison is not the main goal of either experiment.) Other studies find similar results for second-price auctions (Kagel and Levin, 1993) and for ascending clock auctions (McCabe et al., 1990), but these do not directly compare the two formats with the same value distribution and the same subject pool. When we compare **2P** and **AC**, we can see this as a high-powered replication of Kagel et al. (1987), since we now observe 18 groups playing **2P** and 18 groups playing **AC**.³⁶

For the second pair, we compare the second-price plus- X auction (**2P+X**) and the ascending clock plus- X auction (**AC+X**). Subjects' values are drawn as before. However, there is an additional random variable X , which is uniformly distributed between \$0 and \$3. Subjects are not told the value of X until after the auction.

2P+X is SP, but not OSP. In **2P+X**, subjects submit their bids simultaneously. The highest bidder wins the prize if and only if his bid exceeds the second-highest bid plus X . If the highest bidder wins the prize, then he makes a payment equal to the second-highest bid plus X . Otherwise, no agent wins the prize, and no payments are made. In this game, it is a dominant strategy to submit a bid equal to your value.

AC+X is OSP. In **AC+X**, the price starts at a low value (the highest \$25 increment that is below the group draw), and counts upwards. Each bidder can quit at any point.³⁷ When only one bidder is left, the price *continues to rise* for another X dollars, and then

³⁴In both **2P** and **AC**, if there is a tie for the highest bid, then no bidder wins the object.

³⁵During the auction, each bidder observes the number of active bidders.

³⁶I am not aware of any previous laboratory experiment that directly compares second-price and ascending clock auctions, holding constant the value distribution and subject pool, with more than five groups playing each format.

³⁷As in **AC**, each bidder observes the number of active bidders. However, if the number of active bidders is 1 or 2, then the computer display informs bidders that the number of active bidders is "1 or 2".

freezes. If the highest bidder keeps bidding until the price freezes, then she wins the prize at the final price. Otherwise, no agent wins the prize and no payments are made. In this game, it is an obviously dominant strategy to keep bidding if the price is strictly below your value, and quit otherwise.

Some subjects might find **2P** or **AC** familiar, since such mechanisms occur in some natural economic environments. Differences in subject behavior might be caused by different degrees of familiarity with the mechanism. **2P+X** and **AC+X** are novel mechanisms that subjects are unlikely to find familiar. **2P+X** and **AC+X** can be seen as perturbations of **2P** and **AC**; the underlying choice rule is made more complex while preserving the SP-OSP distinction. Thus, comparing **2P+X** and **AC+X** indicates whether the distinction between SP and OSP mechanisms holds for novel and more complicated auction formats.

In the third pair of games, subjects may receive one of four common-value money prizes. The four prize values are drawn, uniformly at random and without replacement, from the set $\{\$0.00, \$0.25, \$0.50, \$0.75, \$1.00, \$1.25\}$. Subjects observe the values of all four prizes at the start of each game.

In a strategy-proof random serial dictatorship (**SP-RSD**), subjects are informed of their priority score, which is drawn uniformly at random from the integers 1 to 10. They then simultaneously submit ranked lists of the four prizes. Players are processed sequentially, from the highest priority score to the lowest. Ties in priority score are broken randomly. Each player is assigned the highest-ranked prize on his list, among the prizes that have not yet been assigned. It is a dominant strategy to rank the prizes in order of their money value. **SP-RSD** is SP, but not OSP.³⁸

In an obviously strategy-proof random serial dictatorship (**OSP-RSD**), subjects are informed of their priority score. Players take turns, from the highest priority score to the lowest. When a player takes his turn, he is shown the prizes that have not yet been taken, and picks one of them. It is an obviously dominant strategy to pick the available prize with the highest money value.

SP-RSD and **OSP-RSD** differ from the auctions in several ways. The auctions are private-value games of incomplete information, whereas **SP-RSD** and **OSP-RSD** are common-value games of complete information. In the auctions, subjects face two sources of strategic uncertainty: They are uncertain about their opponents' valuations, and they are uncertain about their opponents' strategies (a function of valuations). By contrast, in **SP-RSD** and **OSP-RSD**, subjects face no uncertainty about their

³⁸**SP-RSD** is SP, but not OSP since, if a player swaps the order of the highest and second-highest prizes, he might win the second-highest prize. If he reports his true rank-order list, he might win the third-highest prize.

Table 2: Mechanisms in each treatment

	10 rounds	10 rounds	10 rounds
Treatment 1	AC	AC+X	OSP-RSD
Treatment 2	2P	2P+X	SP-RSD
Treatment 3	AC	AC+X	SP-RSD
Treatment 4	2P	2P+X	OSP-RSD

opponents' valuations.

Unlike the auctions, **SP-RSD** and **OSP-RSD** are constant-sum games, such that one player's action cannot affect total player surplus. Any effect that persists in both the auctions and the serial dictatorships is difficult to explain using social preferences, since such theories typically make different predictions for constant-sum and non-constant-sum games. Thus, in comparing **SP-RSD** and **OSP-RSD**, we test whether the SP-OSP distinction has empirical support in mechanisms that are very different from auctions.

At the start of the experiment, subjects are randomly assigned into groups of four. These groups persist throughout the experiment. Consequently, each group's play can be regarded as a single independent observation in the statistical analysis.

Each group either plays 10 rounds of **AC**, followed by 10 rounds of **AC+X**, or plays 10 rounds of **2P**, followed by 10 rounds of **2P+X**.³⁹ At the end of each round, subjects are shown the auction result, their own profit from this round, the winning bidder's profit from this round, and the bids (in order from highest to lowest). Notice that subjects have 10 rounds of experience with a standard auction, before being presented with its unusual **+X** variant. Thus, the data from **+X** auctions record moderately experienced bidders grappling with a new auction format.

Next, groups are re-randomized into either 10 rounds of **OSP-RSD** or 10 rounds of **SP-RSD**. At the end of each round, subjects see which prize they have obtained, and whether their priority score was the highest, or second-highest, and so on.

Table 2 summarizes the design. Subjects had printed copies of the instructions, and the experimenter read aloud the part pertaining to each 10-round segment just before that segment began. The instructions (correctly) informed subjects that their play in earlier segments would not affect the games in later segments. The instructions did not mention dominant strategies or provide recommendations for how to play, so as to prevent confounds from the experimenter demand effect. Instructions for both SP and

³⁹If a stage game with dominant strategies is repeated finitely many times, then the resulting repeated game typically does not have a dominant strategy. The same holds for obviously dominant strategies. Consequently, in interpreting these results as informing us about dominant strategy play, we invoke an implicit narrow framing assumption. The same assumption is made for other experiments in this literature, such as Kagel et al. (1987) and Kagel and Levin (1993).

OSP mechanisms are of similar length and similar reading levels⁴⁰, and can be found in the online appendix.

In every SP mechanism, each subject had 90 seconds to make his choice. Each subject could revise his choice as many times as he desired during the 90 seconds, and only his final choice would count. For OSP mechanisms, mean time to completion was 113.0 seconds in **AC**, 121.4 seconds in **AC+X**, and 40.5 seconds in **OSP-RSD**. However, the rules of the OSP mechanisms imply that not every subject was actively choosing throughout that time.

5.2 Administrative details

Subjects were paid \$20 for participating, in addition to their profits or losses from every round of the experiment. On average, subjects made \$37.54, including the participation payment. Subjects who made negative net profits received just the \$20 participation payment.

I conducted the experiment at the Ohio State University Experimental Economics Laboratory in August 2015, using z-Tree (Fischbacher, 2007). I recruited subjects from the student population using an online system (Greiner, 2015). I administered 16 sessions, where each session involved 1 to 3 groups. Each session lasted about 90 minutes. In total, the data include 144 subjects in 36 groups of 4 (with 9 groups in each treatment).⁴¹

5.3 Statistical Analysis

The data include 4 different auction formats, with 180 auctions per format, for a total of 720 auctions.⁴²

One natural summary statistic for each auction is the difference between the second-highest bid and the second-highest value. This is, equivalently, the difference between that auction's closing price, and the closing price that would have occurred if all bidders played the dominant strategy. Figure 2a displays histograms of the second-highest bid minus the second-highest value, for **AC** and **2P**. Figure 2b does the same for **AC+X** and **2P+X**. If all agents are playing the dominant strategy in an auction, then the histogram for that auction will be a point mass at zero.

⁴⁰Both sets of instructions are approximately at a fifth-grade reading level according to the Flesch-Kincaid readability test, which is a standard measure for how difficult a piece of text is to read (Kincaid et al., 1975).

⁴¹In two cases, network errors caused crashes which prevented a group from continuing in the experiment. I recruited new subjects to replace these groups.

⁴²In 2 out of 720 auctions, computer errors prevented bidders from correctly entering their bids. We omit these 2 observations, but including them has little effect on any of the results that follow.

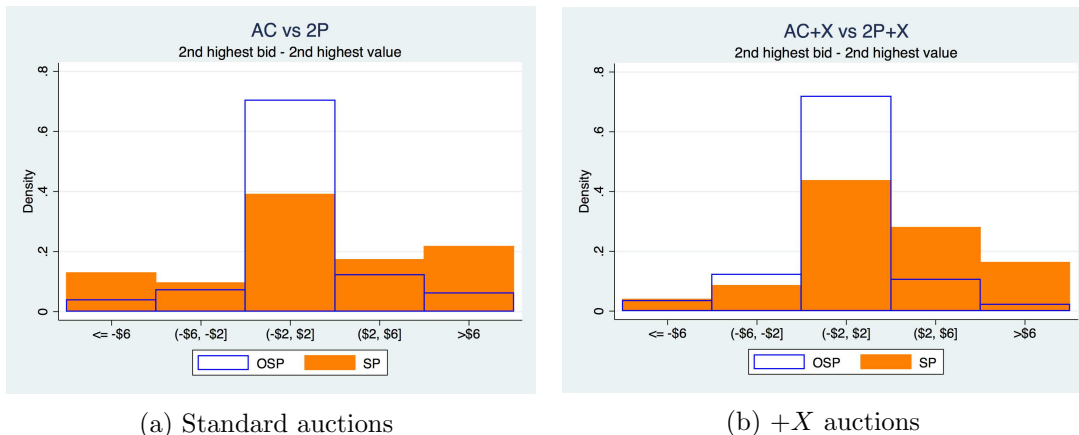


Figure 2: Histogram: 2nd-highest bid minus 2nd-highest value

There is a substantial difference between the empirical distributions for OSP and SP mechanisms. If we choose a random auction from the data, how likely is it to have a closing price within \$2.00 of the equilibrium price? An auction is 31 percentage points more likely to have a closing price within \$2.00 of the equilibrium price under **AC** (OSP) compared to **2P** (SP). An auction is 28 percentage points more likely to have a closing price within \$2.00 of the equilibrium price under **AC+X** (OSP) compared to **2P+X** (SP). Closing prices under **2P+X** are systematically biased upwards ($p = .0031$)⁴³.

Table 3 displays the mean absolute difference between the second-highest bid and the second-highest value, for the first 5 rounds and the last 5 rounds of each auction. This measures the magnitude of errors under each mechanism. (Alternative measures of errors are in Appendix C.) Errors are systematically larger under SP than under OSP, and this difference is significant in both the standard auctions and the novel **+X** auctions, and in both early and late rounds. To build intuition for effect sizes, consider that the expected profit of the winning bidder in **2P** and **AC** is about \$4.00 (given dominant strategy play). Thus, the average errors under **2P** are larger than the theoretical prediction for total bidder surplus.

There is some evidence of learning in **2P**; errors are smaller in the last five rounds compared to the first five rounds ($p = .045$, paired t -test). For the other three auction formats, there is no significant evidence of learning.⁴⁴

To compare subject behavior under **SP-RSD** and **OSP-RSD**, we compute the pro-

⁴³For each group, we take the mean difference between the second-highest bid and the second-highest value. This produces one observation per group playing **2P+X**, for a total of 18 observations, and we use a t -test for the null that these have zero mean.

⁴⁴ $p = .173$ for **AC**, $p = .694$ for **2P+X**, and $p = .290$ for **AC+X**.

Table 3: mean(abs(2nd bid - 2nd value))

Format	Rounds	SP	OSP	p -value
Auction	1-5	8.04	3.19	.006
		(1.25)	(1.05)	
	6-10	4.99	1.77	.016
		(1.18)	(0.33)	
+X Auction	1-5	3.99	1.83	.006
		(0.60)	(0.41)	
	6-10	3.69	1.29	.017
		(0.87)	(0.33)	

For each group, we take the mean absolute difference over each 5-round block. We then compute standard errors counting each group’s 5-round mean as a single observation. (18 observations per cell, standard errors in parentheses.) p -values are computed using a two-sample t -test, allowing for unequal variances. Other empirical strategies yield similar results; see Appendix C for details.

Table 4: Proportion of serial dictatorships not ending in dominant strategy outcome

	SP	OSP	p -value
Rounds 1-5	43.3%	7.8%	.0002
	(7.3%)	(3.3%)	
Rounds 6-10	28.9%	6.7%	.0011
	(5.2%)	(3.2%)	
p -value	.1026	.7492	

For each group, for each 5-round block, we record the error rate. We then compute standard errors counting each group’s observed error rate as a single observation. (18 observations per cell, standard errors in parentheses.) When comparing SP to OSP, we compute p -values using a two-sample t -test, allowing for unequal variances. (Alternative empirical strategies yield similar results. See Appendix C for details.) When comparing early to late rounds of the same game, we compute p -values using a paired t -test.

portion of games that *do not* end in the dominant strategy outcome. Under **SP-RSD**, 36% of games do not end in the dominant strategy outcome. Under **OSP-RSD**, 7% of games do not end in the dominant strategy outcome. Table 4 displays the empirical frequency of non-dominant strategy outcomes, by format and by 5-round blocks. Deviations from the dominant strategy outcome happen more frequently under **SP-RSD** than under **OSP-RSD**, and these differences are highly significant in both early and late rounds.

In **SP-RSD**, 29% of submitted rank-order lists contain errors. The most common error under **SP-RSD** is to swap the ranks of the highest and second-highest prizes, and report the list in order 2nd-1st-3rd-4th. This error accounts for 18% of incorrect rank-order lists.

5.4 One-Shot Treatments

At the suggestion of a referee, I ran additional treatments that investigate **SP-RSD** and **OSP-RSD** without repeated play and with higher stakes, at the same laboratory as the original treatments. These treatments consist of a one-shot serial dictatorship, followed by a demographic questionnaire. The stakes are twelve times as high; four money prizes are drawn uniformly at random and without replacement from the set $\{\$0, \$3, \$6, \$9, \$12, \$15\}$.

I administered 24 sessions, 12 in each treatment, in February 2017. In total, the data include 404 subjects in 101 groups, 48 in one-shot **SP-RSD** and 53 in one-shot **OSP-RSD**.⁴⁵

Under one-shot **SP-RSD**, 40% of games do not end in the dominant strategy outcome, compared to 8% under one-shot **OSP-RSD**. This difference is statistically significant ($p = .0001$, Fisher’s exact test). For comparison, the corresponding error rates in the first round of the ten-round treatments are 44% for **SP-RSD** and 17% for **OSP-RSD**. Error rates in the one-shot treatments are not significantly different from first-round error rates in the ten-round treatments.⁴⁶ In one-shot **SP-RSD**, subjects who submit incorrect lists lose \$2.30 on average, compared to the amount they would have earned by playing the dominant strategy.⁴⁷

5.5 Summary of Experiment

In summary, subjects play the dominant strategy at higher rates in OSP mechanisms, as compared to SP mechanisms that should (according to standard theory) implement the same choice rule. This difference is significant and substantial across all three pairs of mechanisms, and persists after five rounds of repeated play with feedback. Of course, the set of all strategy-proof mechanisms cannot be tested exhaustively in a single experiment. This is only preliminary evidence that obvious dominance correctly classifies strategy-proof mechanisms according to their cognitive complexity.

There are not (yet) well-formed alternative theories that explain the data presented here. The following theories predict dominant-strategy play in any strategy-proof mecha-

⁴⁵The imbalance between treatments is due to subjects who registered to participate, but were absent at the time their session started. Each session lasted about 10 minutes, and subjects were paid \$5 for participating, in addition to their incentive payments.

⁴⁶ $p = .802$ for **SP-RSD** and $p = .176$ for **OSP-RSD**, Fisher’s exact test. Zhang and Levin (2017) investigate a similar strategy-proof serial dictatorship, and find that 37% of games do not end in the dominant-strategy outcome.

⁴⁷In one-shot **SP-RSD**, 33% of subjects submit incorrect rank-order lists. The most common error is still to report the order 2nd-1st-3rd-4th, which accounts for 30% of incorrect lists. Demographic analysis is in Appendix C.

nism: Level- k equilibrium⁴⁸, cognitive hierarchy equilibrium⁴⁹, cursed equilibrium⁵⁰, and analogy-based expectation equilibrium⁵¹. Standard specifications of quantal response equilibrium⁵² predict substantial under-bidding in ascending auctions, and simulations suggest that it predicts *larger* errors in ascending auctions than in second-price auctions, which is the opposite of the effect in the data (Appendix D provides details).

6 Discussion

Sometimes, it is important to design simple mechanisms. If simple auctions attract more participants, then they can raise more revenue than complex optimal auctions (Bulow and Klemperer, 1996). Moreover, human beings tend to play their equilibrium strategies in simple mechanisms, but make large mistakes in complex mechanisms.⁵³ This affects the performance of complex optimal mechanisms, often in ways that are difficult to foresee. Simplicity can be a real constraint - every bit as real as the constraints imposed by technology or incentives.⁵⁴

Consequently, we need a precise definition of simplicity. Of course, mechanisms can be simple or complex in many ways, and the definition I have provided captures only one of them. Any definition must trade off multiple worthy concerns. It must be empirically accurate, but analytically tractable. It must be weak enough to allow positive results, but strong enough to discipline our choice of mechanisms.

Here are some dimensions of simplicity that OSP does not capture, which might be fruitful for investigation: Firstly, since OSP invokes only standard game-theoretic primitives, it does not capture framing effects. Any multiple-choice mathematics quiz has an obviously dominant strategy, since there is a deterministic correct answer for every question. We can embed puzzles of arbitrary complexity into the labels of moves, which OSP entirely neglects. Secondly, OSP abstracts from the complexity of searching a deep game tree for the optimal strategy. The agent compares his obviously dominant strategy to deviations at each information set, but how does he find that strategy in the first place? In a well-engineered mechanism, the answer may be that the designer draws the agent's attention to it.

⁴⁸Stahl and Wilson (1994, 1995); Nagel (1995).

⁴⁹Camerer et al. (2004).

⁵⁰Eyster and Rabin (2005); Esponda (2008).

⁵¹Jehiel (2005).

⁵²McKelvey and Palfrey (1995, 1998)

⁵³Such behavior has been found in the laboratory (Kagel et al., 1987; Kagel and Levin, 1993; Chen and Sönmez, 2006; Zhang and Levin, 2017), and in high-stakes decisions in the field (Hassidim et al., 2016; Rees-Jones, 2017).

⁵⁴This paraphrases an observation that Roth (2007) makes about repugnance.

A formal definition lets us quantify the cost of simplicity. For instance, it is complex to sell multiple objects in a combinatorial auction, because the welfare-optimal mechanism must elicit preferences over the power set of objects. In such environments, it can be that no OSP mechanism yields first-best welfare.⁵⁵ However, when objects are substitutes and types are independently distributed, there always exists an OSP mechanism that delivers at least half of first-best expected welfare (Feldman et al., 2014; Dütting et al., 2016).⁵⁶ In these environments, the ‘cost of simplicity’ is bounded by a half.

When designing mechanisms, the appropriate constraints depend on the intended application. (The requirement that a ship withstand thirty atmospheres of pressure is superfluous for many vessels, but vital for submarines.) If the agents are expert strategists, and the mechanism will be thoroughly audited by a trusted third-party, then OSP is too demanding. Strategically sophisticated agents and full commitment power are the ideal conditions for mechanism design. These conditions do not always obtain, so it is valuable to design mechanisms that are robust to weaker assumptions.

References

- Ashlagi, I. and Gonczarowski, Y. A. (2015). No stable matching mechanism is obviously strategy-proof. *arXiv preprint arXiv:1511.00452*.
- Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *American Economic Review*, pages 1452–1475.
- Bade, S. and Gonczarowski, Y. A. (2016). Gibbard-Satterthwaite success stories and obvious strategyproofness. *working paper*.
- Barberà, S., Berga, D., and Moreno, B. (2016). Group strategy-proofness in private good economies. *American Economic Review*, 106(4):1073–99.
- Bartal, Y., Gonen, R., and Nisan, N. (2003). Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 72–87. ACM.
- Bird, C. G. (1984). Group incentive compatibility in a market with indivisible goods. *Economics Letters*, 14(4):309–313.

⁵⁵This holds even when objects have additively separable values (Bade and Gonczarowski, 2016).

⁵⁶This statement is entailed by Lemma 3.4 of Feldman et al. (2014). It even holds for the more general class of fractionally subadditive preferences.

- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, pages 180–194.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, pages 861–898.
- Charness, G. and Levin, D. (2009). The origin of the winner’s curse: a laboratory study. *American Economic Journal: Microeconomics*, pages 207–236.
- Chen, Y. and Sönmez, T. (2006). School choice: an experimental study. *Journal of Economic theory*, 127(1):202–231.
- Cramton, P. (1998). Ascending auctions. *European Economic Review*, 42(3):745–756.
- Dütting, P., Feldman, M., Kesselheim, T., and Lucier, B. (2016). Posted prices, smoothness, and combinatorial prophet inequalities. *arXiv preprint arXiv:1612.03161*.
- Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, pages 242–259.
- Elmes, S. and Reny, P. J. (1994). On the strategic equivalence of extensive form games. *Journal of Economic Theory*, 62(1):1–23.
- Esponda, I. (2008). Behavioral equilibrium in economies with adverse selection. *The American Economic Review*, 98(4):1269–1291.
- Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202.
- Esponda, I. and Vespa, E. (2016). Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. *working paper*.
- Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.
- Feldman, M., Gravin, N., and Lucier, B. (2014). Combinatorial auctions via posted prices. *CoRR*, abs/1411.4916.
- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Friedman, E. J. (2002). Strategic properties of heterogeneous serial cost sharing. *Mathematical Social Sciences*, 44(2):145–154.

- Friedman, E. J. (2004). Asynchronous learning in decentralized environments: A game-theoretic approach. In *Collectives and the Design of Complex Systems*, pages 133–143. Springer.
- Glazer, J. and Rubinstein, A. (1996). An extensive game as a guide for solving a normal game. *Journal of Economic Theory*, 1(70):32–42.
- Green, J. and Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, pages 427–438.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125.
- Harstad, R. M. (2000). Dominant strategy adoption and bidders’ experience with pricing rules. *Experimental economics*, 3(3):261–280.
- Hassidim, A., Marciano-Romm, D., Romm, A., and Shorrer, R. I. (2016). “Strategic” behavior in a strategy-proof environment. *working paper*.
- Holmström, B. (1979). Groves’ scheme on restricted domains. *Econometrica*, pages 1137–1144.
- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic theory*, 123(2):81–104.
- Kagel, J. H., Harstad, R. M., and Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):pp. 1275–1304.
- Kagel, J. H. and Levin, D. (1993). Independent private value auctions: Bidder behaviour in first-, second- and third-price auctions with varying numbers of bidders. *The Economic Journal*, 103(419):868–879.
- Kincaid, J. P., Fishburne Jr, R. P., Rogers, R. L., and Chissom, B. S. (1975). Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, DTIC Document.
- Li, S. (2017). Obvious ex post equilibrium. *working paper*.
- Loertscher, S. and Marx, L. M. (2017). Optimal clock auctions. *working paper*.

- McCabe, K. A., Rassenti, S. J., and Smith, V. L. (1990). Auction institutional design: Theory and behavior of simultaneous multiple-unit generalizations of the Dutch and English auctions. *The American Economic Review*, pages 1276–1283.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38.
- McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental economics*, 1(1):9–41.
- Milgrom, P. and Segal, I. (2015). Deferred-acceptance auctions and radio spectrum reallocation. *working paper*.
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, pages 175–208.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, pages 1313–1326.
- Ngangoue, K. and Weizsacker, G. (2015). Learning from unrealized versus realized prices. *working paper*.
- Pathak, P. A. and Sönmez, T. (2008). Leveling the playing field: Sincere and sophisticated players in the Boston mechanism. *The American Economic Review*, 98(4):1636–1652.
- Pycia, M. and Troyan, P. (2016). Obvious dominance and random priority. *working paper*.
- Rees-Jones, A. (2017). Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match. *SSRN working paper*.
- Roth, A. E. (1982). Incentive compatibility in a market with indivisible goods. *Economics Letters*, 9(2):127–132.
- Roth, A. E. (2007). Repugnance as a constraint on markets. *Journal of Economic Perspectives*, 21(3):37–58.
- Rothkopf, M. H., Teisberg, T. J., and Kahn, E. P. (1990). Why are Vickrey auctions rare? *Journal of Political Economy*, pages 94–109.

- Saks, M. and Yu, L. (2005). Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, pages 286–293. ACM.
- Savage, L. J. (1954). *The Foundations of Statistics*. Wiley & Sons, New York.
- Shapley, L. and Scarf, H. (1974). On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37.
- Shimoji, M. and Watson, J. (1998). Conditional dominance, rationalizability, and game forms. *Journal of Economic Theory*, 83(2):161–195.
- Spence, M. (1974). Competitive and optimal responses to signals: An analysis of efficiency and distribution. *Journal of Economic Theory*, 7(3):296–332.
- Stahl, D. O. and Wilson, P. W. (1994). Experimental evidence on players’ models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327.
- Stahl, D. O. and Wilson, P. W. (1995). On players’ models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.
- Thompson, F. (1952). Equivalence of games in extensive form, RM 759. *RAND Corporation*.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.
- Zhang, L. and Levin, D. (2017). Partition obvious preference and mechanism design: Theory and experiment. *working paper*.

A Definition of Extensive Game Forms with Consequences in X

An *extensive game form with consequences in X* is a tuple $\langle H, \prec, A, \mathcal{A}, P, \delta_c, (\mathcal{I}_i)_{i \in N}, g \rangle$, where:

1. H is a set of histories, along with a binary relation \prec on H that represents precedence.
 - (a) \prec is a partial order, and (H, \prec) form an arborescence⁵⁷.

⁵⁷That is, a directed rooted tree such that every edge points away from the root.

(b) h_\emptyset denotes $h \in H : \neg \exists h' : h' \prec h$

(c) H has finite depth, i.e.:

$$\exists k \in \mathbb{N} : \forall h \in H : |\{h' \in H : h' \prec h\}| \leq k \quad (14)$$

(d) $Z \equiv \{h \in H : \neg \exists h' : h \prec h'\}$

(e) $\sigma(h)$ denotes the set of immediate successors of h .

2. A is a set of actions.

3. $\mathcal{A} : H \setminus h_\emptyset \rightarrow A$ labels each non-initial history with the last action taken to reach it.

(a) For all h , \mathcal{A} is one-to-one on $\sigma(h)$.

(b) $A(h)$ denotes the actions available at h .

$$A(h) \equiv \bigcup_{h' \in \sigma(h)} \mathcal{A}(h') \quad (15)$$

4. P is a player function. $P : H \setminus Z \rightarrow N \cup c$, where c is chance.

5. δ_c is the chance function. It specifies a probability measure over chance moves. d_c denotes some realization of chance moves: For any h where $P(h) = c$, $d_c(h) \in A(h)$.⁵⁸

6. \mathcal{I}_i is a partition of $\{h : P(h) = i\}$ such that:

(a) $A(h) = A(h')$ whenever h and h' are in the same cell of the partition.

(b) For any $I_i \in \mathcal{I}_i$, we denote: $P(I_i) \equiv P(h)$ for any $h \in I_i$. $A(I_i) \equiv A(h)$ for any $h \in I_i$.

(c) Each action is available at only one information set: If $a \in A(I_i)$, $a' \in A(I'_j)$, $I_i \neq I'_j$ then $a \neq a'$.

7. g is an outcome function. It associates each terminal history with an outcome.

$$g : Z \rightarrow X$$

⁵⁸We could make the addition assumption that δ_c has full support on the available moves $A(h)$ when it is called to play. This ensures a pleasing invariance property: It rules out games with zero-probability chance moves that do not affect play, but do affect whether a strategy is obviously dominant. However, a full support assumption is not necessary for any of the results that follow, so we do not make it here.

B Proofs omitted from the main text (for online publication)

B.1 Theorem 1

Proof. First we prove the “if” direction. Fix agent 1 and preferences θ_1 . Suppose that S_1 is not obviously dominant in $G = \langle H, \prec, A, \mathcal{A}, P, \delta_c, (\mathcal{I}_i)_{i \in N}, g \rangle$. We need to demonstrate that there exists \tilde{G} that is i -indistinguishable from G , such that $\lambda_{G, \tilde{G}}(S_1)$ is not weakly dominant in \tilde{G} . We proceed by construction.

Let $(S'_1, I_1, h^{\text{sup}}, S_{-1}^{\text{sup}}, d_c^{\text{sup}}, h^{\text{inf}}, S_{-1}^{\text{inf}}, d_c^{\text{inf}})$ be such that $I_1 \in \alpha(S_1, S'_1)$, $h^{\text{inf}} \in I_1$, $h^{\text{sup}} \in I_1$, and

$$u_1^G(h^{\text{sup}}, S'_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}}, \theta_1) > u_1^G(h^{\text{inf}}, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}}, \theta_1) \quad (16)$$

Since G is a game of perfect recall, we can pick $(S_{-1}^{\text{inf}}, d_c^{\text{inf}})$ such that $h^{\text{inf}} \prec z^G(h_\emptyset, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}})$, by specifying that $(S_{-1}^{\text{inf}}, d_c^{\text{inf}})$ plays in a way consistent with h^{inf} at any $h \prec h^{\text{inf}}$. Likewise for h^{sup} and $(S_{-1}^{\text{sup}}, d_c^{\text{sup}})$. Suppose we have so done.

We now define another game \tilde{G} is 1-indistinguishable from G . Intuitively, we construct this as follows:

1. We add a chance move at the start of the game; chance can play L or R .
2. Agent 1 does not at any history know whether chance played L or R .
3. If chance plays L , then the game proceeds as in G .
4. If chance plays R , then the game proceeds mechanically as though all players in $N \setminus 1$ and chance played according to $S_{-1}^{\text{sup}}, d_c^{\text{sup}}$ in G , with one exception:
5. If chance played R , we reach the information set corresponding to I_1 , and agent 1 plays $S_1(I_1)$, then the game henceforth proceeds mechanically as though all players in $N \setminus 1$ and chance played according to $S_{-1}^{\text{inf}}, d_c^{\text{inf}}$ in G .

Formally, the construction proceeds as such: $\tilde{A} = A \cup \{L, R\}$, where $A \cap \{L, R\} = \emptyset$. There is a new starting history \tilde{h}_\emptyset , with two successors $\sigma(\tilde{h}_\emptyset) = \{\tilde{h}_L, \tilde{h}_R\}$, $\tilde{\mathcal{A}}(\tilde{h}_L) = L$, $\tilde{\mathcal{A}}(\tilde{h}_R) = R$, $\tilde{P}(\tilde{h}_\emptyset) = c$. The subtree $\tilde{H}_L \subset \tilde{H}$ starting from \tilde{h}_L ordered by \tilde{z} is the same as the arborescence (H, \prec) . $(\tilde{\mathcal{A}}, \tilde{P}, \tilde{\delta}_c, \tilde{g})$ are defined on \tilde{H}_L exactly as $(\mathcal{A}, P, \delta_c, g)$ are on H . For $j \neq 1$, $\tilde{\mathcal{I}}_j$ is defined as on H .

We now construct the subtree starting from \tilde{h}_R . Let h^* be such that $h^* \in \sigma(h^{\text{sup}})$, $h^* \preceq z^G(h^{\text{sup}}, S_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}})$.

$$\begin{aligned}
H' \equiv & \{h \in H \mid \exists S_1'' : h \preceq z^G(h_\emptyset, S_1'', S_{-1}^{\text{sup}}, d_c^{\text{sup}})\} \\
& \cap [\{h \in H \mid P(h) = 1\} \cup \{h \in Z\}] \\
& \setminus \{h \in H \mid h^* \preceq h\}
\end{aligned} \tag{17}$$

In words, these are the histories that can be reached by *some* S_1'' when facing $S_{-1}^{\text{sup}}, d_c^{\text{sup}}$, where either agent 1 is called to play or that history is terminal, and such that those histories are not h^* or its successors.

Let h^{**} be such that $h^{**} \in \sigma(h^{\text{inf}})$, $h^{**} \preceq z^G(h^{\text{inf}}, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}})$.

$$\begin{aligned}
H'' \equiv & \{h \in H \mid \exists S_1'' : h \preceq z^G(h_\emptyset, S_1'', S_{-1}^{\text{inf}}, d_c^{\text{inf}})\} \\
& \cap [\{h \in H \mid P(h) = 1\} \cup \{h \in Z\}] \\
& \cap \{h \in H \mid h^{**} \preceq h\}
\end{aligned} \tag{18}$$

In words, these are the histories that can be reached by *some* S_1'' when facing $S_{-1}^{\text{inf}}, d_c^{\text{inf}}$, where either agent 1 is called to play or that history is terminal, and such that those histories are h^{**} or its successors.

We now paste these together. Let \tilde{H}_R be the rooted subtree ordered by $\tilde{\prec}$, for some bijection $\gamma : \tilde{H}_R \rightarrow H' \cup H''$, such that for all $\tilde{h}, \tilde{h}' \in \tilde{H}_R$, $\tilde{h} \tilde{\prec} \tilde{h}'$ if and only if

1. EITHER: $\gamma(\tilde{h}), \gamma(\tilde{h}') \in H'$ and $\gamma(\tilde{h}) \prec \gamma(\tilde{h}')$
2. OR: $\gamma(\tilde{h}), \gamma(\tilde{h}') \in H''$ and $\gamma(\tilde{h}) \prec \gamma(\tilde{h}')$
3. OR: $\gamma(\tilde{h}) \prec h^*$ and $h^{**} \preceq \gamma(\tilde{h}')$

The root of this subtree exists and is unique; it corresponds to $\gamma^{-1}(h)$, where h is the earliest history preceding $z^G(h_\emptyset, S_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}})$ where 1 is called to play. Let \tilde{h}_R be the root of \tilde{H}_R . This completes the specification of \tilde{H} .

For all $\tilde{h} \in \tilde{H}_R$, we define:

1. $\tilde{g}(\tilde{h}) = g(\gamma(\tilde{h}))$ if \tilde{h} is a terminal history.
2. $\tilde{P}(\tilde{h}) = 1$ if \tilde{h} is not a terminal history.

For all $\tilde{h} \in \tilde{H}_R \setminus \tilde{h}_R$, we define $\tilde{A}(\tilde{h}) = A(h)$, for the unique (\tilde{h}', h) such that:

1. $\tilde{h} \in \sigma(\tilde{h}')$

2. $h \in \sigma(\gamma(\tilde{h}'))$

3. $h \preceq \gamma(\tilde{h})$

We now specify the information sets for agent 1. Every $\tilde{h} \in \tilde{H}_L$ corresponds to a unique history in H . We use γ_L to denote the bijection from \tilde{H}_L to H . Let $\hat{\gamma}$ be defined as γ_L on \tilde{H}_L and γ on \tilde{H}_R .

1's information partition $\tilde{\mathcal{I}}_1$ is defined as such: $\forall \tilde{h}, \tilde{h}' \in \tilde{H}$:

$$\exists \tilde{I}'_1 \in \tilde{\mathcal{I}}_1 : \tilde{h}, \tilde{h}' \in \tilde{I}'_1 \quad (19)$$

if and only if

$$\exists I'_1 \in \mathcal{I}_1 : \hat{\gamma}(\tilde{h}), \hat{\gamma}(\tilde{h}') \in I'_1 \quad (20)$$

All that remains is to define δ_c ; we need only specify that at \tilde{h}_\emptyset , c plays R with certainty.⁵⁹

$\tilde{G} = \langle \tilde{H}, \tilde{z}, \tilde{A}, \tilde{\mathcal{A}}, \tilde{P}, \tilde{\delta}_c, (\tilde{\mathcal{I}}_i)_{i \in N}, \tilde{g} \rangle$ is 1-indistinguishable from G . Every experience at some history in \tilde{H}_L corresponds to some experience in G , and *vice versa*. Moreover, any experience at some history in \tilde{H}_R could also be produced by some history in \tilde{H}_L .

Let $\lambda_{G, \tilde{G}}$ be the appropriate bijection from 1's information sets and actions in G onto 1's information sets and actions in \tilde{G} . Take arbitrary \tilde{S}_{-1} . Observe that since $I_1 \in \alpha(S_1, S'_1)$, $\lambda_{G, \tilde{G}}(S_1)$ and $\lambda_{G, \tilde{G}}(S'_1)$ result in the same histories following \tilde{h}_R , until they reach information set $\lambda_{G, \tilde{G}}(I_1)$. Having reached that point, $\lambda_{G, \tilde{G}}(S_1)$ leads to outcome $g(z^G(h^{\text{inf}}, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}}))$ and $\lambda_{G, \tilde{G}}(S'_1)$ leads to outcome $g(z^G(h^{\text{sup}}, S'_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}}))$. Thus,

$$\begin{aligned} & \mathbb{E}_{\tilde{\delta}_c} [u_1^{\tilde{G}}(\tilde{h}_\emptyset, \lambda_{G, \tilde{G}}(S'_1), \tilde{S}_{-1}, \tilde{d}_c, \theta_1)] \\ &= u_1^G(h^{\text{sup}}, S'_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}}, \theta_1) \\ &> u_1^G(h^{\text{inf}}, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}}, \theta_1) \\ &= \mathbb{E}_{\tilde{\delta}_c} [u_1^{\tilde{G}}(\tilde{h}_\emptyset, \lambda_{G, \tilde{G}}(S_1), \tilde{S}_{-1}, \tilde{d}_c, \theta_1)] \end{aligned} \quad (21)$$

So $\lambda_{G, \tilde{G}}(S_1)$ is not weakly dominant in \tilde{G} .

We now prove the “only if” direction. Take arbitrary \tilde{G} . Suppose $\lambda_{G, \tilde{G}}(S_1) \equiv \tilde{S}_1$ is not weakly dominant in \tilde{G} (given type θ_1). We want to show that S_1 is not obviously dominant in G .

⁵⁹If one prefers to avoid δ_c without full support, an alternative proof puts ϵ probability on L , for ϵ close to 0, or for games with $|N| > 2$ assigns $\tilde{P}(\tilde{h}_\emptyset) = 2$.

There exist \tilde{S}'_1 and \tilde{S}'_{-1} such that:

$$\mathbb{E}_{\tilde{\delta}_c}[u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1)] > \mathbb{E}_{\tilde{\delta}_c}[u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1)] \quad (22)$$

This inequality must hold for some realization of the chance function, so there exists \tilde{d}_c such that:

$$u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1) > u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1) \quad (23)$$

Fix $(\tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1)$.

$$z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c) \neq z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c) \quad (24)$$

Define:

$$\tilde{H}^* \equiv \{\tilde{h} \in \tilde{H} \mid \tilde{h} \prec z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c) \text{ and } \tilde{h} \prec z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c)\} \quad (25)$$

$$\tilde{h}^* \equiv \tilde{h} \in \tilde{H}^* \text{ such that } \forall \tilde{h}' \in \tilde{H}^* : \tilde{h}' \preceq \tilde{h} \quad (26)$$

Since the opponent strategies and chance moves are held constant across both sides of Equation 24, $P(\tilde{h}^*) = 1$ and $\tilde{h}^* \in \tilde{I}_1$, where $\tilde{S}_1(\tilde{I}_1) \neq \tilde{S}'_1(\tilde{I}_1)$. Moreover, $\tilde{I}_1 \in \alpha(\tilde{S}_1, \tilde{S}'_1)$ and $\lambda_{\tilde{G}, G}(\tilde{I}_1) \in \alpha(S_1, S'_1)$, where we denote $S'_1 \equiv \lambda_{\tilde{G}, G}(\tilde{S}'_1)$.

Since G and \tilde{G} are 1-indistinguishable, consider the experiences $\lambda_{\tilde{G}, G}(\psi_1(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c)))$ and $\lambda_{\tilde{G}, G}(\psi_1(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c)))$.

In G , $\lambda_{\tilde{G}, G}(\psi_1(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c)))$ could lead to outcome $\tilde{g}(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c))$. We use $(S_{-1}^{\text{inf}}, d_c^{\text{inf}})$ to denote the corresponding opponent strategies and chance realizations that lead to that outcome. We denote $h^{\text{inf}} \equiv h \in \lambda_{\tilde{G}, G}(\tilde{I}_1) : h \prec z^G(h_\emptyset, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}})$.

In G , $\lambda_{\tilde{G}, G}(\psi_1(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c)))$ could lead to outcome $\tilde{g}(z^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c))$. We use $(S_{-1}^{\text{sup}}, d_c^{\text{sup}})$ to denote the corresponding opponent strategies and chance realizations that lead to that outcome. We denote $h^{\text{sup}} \equiv h \in \lambda_{\tilde{G}, G}(\tilde{I}_1) : h \prec z^G(h_\emptyset, S'_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}})$.

$$\begin{aligned} & u_1^G(h^{\text{sup}}, S'_1, S_{-1}^{\text{sup}}, d_c^{\text{sup}}, \theta_1) \\ &= u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}'_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1) \\ &> u_1^{\tilde{G}}(\tilde{h}_\emptyset, \tilde{S}_1, \tilde{S}'_{-1}, \tilde{d}_c, \theta_1) \\ &= u_1^G(h^{\text{inf}}, S_1, S_{-1}^{\text{inf}}, d_c^{\text{inf}}, \theta_1) \end{aligned} \quad (27)$$

where $h^{\text{sup}}, h^{\text{inf}} \in \lambda_{\tilde{G}, G}(\tilde{I}_1)$ and $\lambda_{\tilde{G}, G}(\tilde{I}_1) \in \alpha(S_1, S'_1)$. Thus S_1 is not obviously dominant in G . □

B.2 Theorem 2

Proof. The key is to see that, for every $G \in \mathcal{G}$, there is a corresponding \tilde{S}_0^Δ , and *vice versa*. We use \bar{S}_0 to denote the support of \tilde{S}_0^Δ . In particular, observe the following isomorphism:

Table 5: Equivalence between extensive game forms and Planner mixed strategies

G	\tilde{S}_0^Δ
d_c	$\tilde{S}_0 \in \bar{S}_0$
δ_c	the probability measure specified by \tilde{S}_0^Δ
$g(z)$ for $z \in Z$	the Planner's choice of outcome when she ends the game
I_i	$((m_k, R_k, r_k)_{k=1}^{t-1}, m_t, R_t)$ consistent with some $\tilde{S}_0 \in \bar{S}_0$
$A(I_i)$	R_t
$\psi_i(z)$	o_i^c consistent with some $\tilde{S}_0 \in \bar{S}_0$ and \tilde{S}_N

Information sets in G are equivalent to sequences of past communication $((m_k, R_k, r_k)_{k=1}^{t-1}, m_t, R_t)$ under \tilde{S}_0^Δ . Available actions at some information set $A(I_i)$ are equivalent to acceptable responses R_t . Thus, for any strategy in some game G , we can construct an equivalent strategy given appropriate \tilde{S}_0^Δ , and vice versa.

Furthermore, fixing a chance realization d_c and agent strategies S_N uniquely results in some outcome. Similarly, fixing a realization of the planner's mixed strategy $\tilde{S}_0 \in \bar{S}_0$ and agent strategies \tilde{S}_N uniquely determines some outcome. Consequently, for any $G \in \mathcal{G}$, there exists \tilde{S}_0^Δ with the same strategies available for each agent and the same resulting (probability measure over) outcomes, and *vice versa*.⁶⁰ Table 5 summarizes.

The next step is to see that a bilateral commitment \hat{S}_0^i is equivalent to the Planner promising to 'run' only games in some equivalence class that is i -indistinguishable.

Suppose that there is some G that OSP-implements f . Pick some equivalent \tilde{S}_0^Δ with support \bar{S}_0 . For each $i \in N$, specify the bilateral commitment $\hat{S}_0^i \equiv \Phi_i^{-1}(\Phi_i(\bar{S}_0))$. These bilateral commitments support f .

To see this, take any $\tilde{S}_0^{\Delta'} \in \Delta \hat{S}_0^i$, with support \bar{S}'_0 . For any $\tilde{S}'_0 \in \bar{S}'_0$, for any \tilde{S}'_N , there exists $\tilde{S}_0 \in \bar{S}_0$ and \tilde{S}_N such that $\phi_i(\tilde{S}'_0, \tilde{S}'_N) = \phi_i(\tilde{S}_0, \tilde{S}_N)$. By construction, G is such that: There exists $z \in Z$ where $\psi_i(z)$ and $g(z)$ are equivalent to $\phi_i(\tilde{S}_0, \tilde{S}_N)$. Thus, for

⁶⁰Implicitly, this relies on the requirement that both G and \tilde{S}_0^Δ have finite length. If one had finite length but the other could be infinitely long, the resulting outcome would not be well defined and the equivalence would not hold.

G' that is equivalent to $\tilde{S}_0^{\Delta'}$, every terminal history in G' results in the same experience for i and the same outcome as some terminal history in G . Consequently, G and G' are i -indistinguishable. Thus, by Theorem 1, the strategy assigned to agent i with type θ_i is weakly dominant in G' , which implies that it is a best response to $\tilde{S}_0^{\Delta'}$ and any $\tilde{S}_{N \setminus i}$ in the bilateral commitment game. Thus, if f is OSP-implementable, then f can be supported by bilateral commitments.

Suppose that f can be supported by bilateral commitments $(\hat{S}_0^i)_{i \in N}$, with requisite \tilde{S}_0^{Δ} (with support \bar{S}_0) and $\tilde{\mathbf{S}}_N$. Without loss of generality, let us suppose these are ‘minimal’ bilateral commitments, *i.e.* $\hat{S}_0^i = \Phi_i^{-1}(\Phi_i(\bar{S}_0))$. Pick G that is equivalent to \tilde{S}_0^{Δ} . G OSP-implements f .

To see this, consider any G' such that G and G' are i -indistinguishable. Let $\tilde{S}_0^{\Delta'}$ denote the Planner strategy that corresponds to G' . At any terminal history z' in G' , the resulting experience $\psi_i(z')$ and outcome $g'(z')$ are equivalent to the experience $\psi_i(z)$ and outcome $g(z)$ for some terminal history z in G . These in turn correspond to some observation $o_i \in \Phi_i(\bar{S}_0)$. Thus $\tilde{S}_0^{\Delta'} \in \Delta \hat{S}_0^i$. Since f is supported by $(\hat{S}_0^j)_{j \in N}$, $\tilde{\mathbf{S}}_i^{\theta_i}$ is a best response (for type θ_i) to $\tilde{S}_0^{\Delta'}$ and any $\tilde{S}_{N \setminus i}$. Thus, the equivalent strategy $\mathbf{S}_i(\theta_i)$ is weakly dominant in G' . Since this argument holds for all i -indistinguishable G' , by Theorem 1, $\mathbf{S}_i(\theta_i)$ is obviously dominant in G . Thus, if f can be supported by bilateral commitments, then f is OSP-implementable. □

B.3 Proposition 2

Proof. We prove the contrapositive. Suppose $(\tilde{G}, \tilde{\mathbf{S}}_N)$ does not OSP-implement f . Then there exists some $(i, \theta_i, \tilde{S}'_i, \tilde{I}_i)$ such that $\tilde{I}_i \in \alpha(\tilde{\mathbf{S}}_i(\theta_i), \tilde{S}'_i)$ and

$$u_i^{\tilde{G}}(\tilde{h}, \tilde{\mathbf{S}}_i(\theta_i), \tilde{S}_{-i}, \tilde{d}_c, \theta_i,) < u_i^{\tilde{G}}(\tilde{h}', \tilde{S}'_i, \tilde{S}'_{-i}, \tilde{d}'_c, \theta_i,) \quad (28)$$

for some $(\tilde{h}, \tilde{S}_{-i}, \tilde{d}_c)$ and $(\tilde{h}', \tilde{S}'_{-i}, \tilde{d}'_c)$.

Notice that \tilde{h} and \tilde{h}' correspond to histories h and h' in G . Moreover, we can define $S'_i = \tilde{S}'_i$ at information sets containing histories that are shared by G and \tilde{G} , and specify S'_i arbitrarily elsewhere. We do the same for $(\tilde{S}_{-i}, \tilde{d}_c)$ and $(\tilde{S}'_{-i}, \tilde{d}'_c)$, to construct (S_{-i}, d_c) and (S'_{-i}, d'_c) . But, starting from h and h' respectively, these result in the same outcomes as their partners in \tilde{G} . Thus,

$$u_i^G(h, \mathbf{S}_i(\theta_i), S_{-i}, d_c, \theta_i,) < u_i^G(h', S'_i, S'_{-i}, d'_c, \theta_i,) \quad (29)$$

$h, h' \in I_i$, for $I_i \in \alpha(\mathbf{S}_i(\theta_i), S'_i)$. Thus (G, \mathbf{S}_N) does not OSP-implement f .

□

B.4 Theorem 3

We split this proof into two parts.

Proposition 6. *If (G, \mathbf{S}_N) OSP-implements f_y , then $\mathcal{P}(G, \mathbf{S}_N)$ is a personal-clock auction.*

Proof. Throughout this proof, we use the following notation: Given some type-strategy \mathbf{S}_i , $S_i^{\theta_i}$ denotes the strategy assigned to θ_i by \mathbf{S}_i .

Take any (G, \mathbf{S}_N) that implements (f_y, f_t) . For any history h , we define

$$\Theta_h \equiv \{\theta_N \mid h \preceq z^G(h_\emptyset, (S_i^{\theta_i})_{i \in N})\} \quad (30)$$

$$\Theta_{h,i} \equiv \{\theta_i \mid \exists \theta_{-i} : (\theta_i, \theta_{-i}) \in \Theta_h\} \quad (31)$$

For information set I_i , we define

$$\Theta_{I_i} \equiv \cup_{h \in I_i} \Theta_h \quad (32)$$

$$\Theta_{I_i,i} \equiv \{\theta_i \mid \exists \theta_{-i} : (\theta_i, \theta_{-i}) \in \Theta_{I_i}\} \quad (33)$$

$$\Theta_{I_i,i}^1 \equiv \{\theta_i \mid \exists \theta_{-i} : (\theta_i, \theta_{-i}) \in \Theta_{I_i} \text{ and } i \in f_y(\theta_i, \theta_{-i})\} \quad (34)$$

$$\Theta_{I_i,i}^0 \equiv \{\theta_i \mid \exists \theta_{-i} : (\theta_i, \theta_{-i}) \in \Theta_{I_i} \text{ and } i \notin f_y(\theta_i, \theta_{-i})\} \quad (35)$$

Some observations about this construction:

1. Since player i 's type-strategy depends only on his own type, $\Theta_{I_i,i} = \Theta_{h,i}$ for all $h \in I_i$.
2. $\Theta_{I_i,i} = \Theta_{I_i,i}^1 \cup \Theta_{I_i,i}^0$
3. Since SP requires $1_{i \in f_y(\theta)}$ weakly increasing in θ_i , $\Theta_{I_i,i}^1$ dominates $\Theta_{I_i,i}^0$ in the strong set order.

Lemma 1. *Suppose (G, \mathbf{S}_N) OSP-implements (f_y, f_t) , where $G \equiv \langle H, \prec, A, \mathcal{A}, P, \delta_c, (\mathcal{I}_i)_{i \in N}, g \rangle$. For all i , for all $I_i \in \mathcal{I}_i$, if:*

1. $\theta_i < \theta'_i$
2. $\theta_i \in \Theta_{I_i, i}^1$
3. $\theta'_i \in \Theta_{I_i, i}^0$

then $S_i^{\theta_i}(I_i) = S_i^{\theta'_i}(I_i)$.

Equivalently, for any I_i , there exists $a_{I_i}^*$ such that for all $\theta_i \in \Theta_{I_i, i}^1 \cap \Theta_{I_i, i}^0$, $S_i^{\theta_i}(I_i) = a_{I_i}^*$.

Suppose not. Take $(i, I_i, \theta_i, \theta'_i)$ constituting a counterexample to Lemma 1. Since $\theta_i \in \Theta_{I_i, i}^1$, there exists $h \in I_i$ and S_{-i} such that $i \in g_y(z^G(h, S_i^{\theta_i}, S_{-i}))$. Fix $t_i \equiv g_{t, i}(z^G(h, S_i^{\theta_i}, S_{-i}))$. Since $\theta'_i \in \Theta_{I_i, i}^0$, there exists $h' \in I_i$ and S'_{-i} such that $i \notin g_y(z^G(h', S_i^{\theta'_i}, S'_{-i}))$. Fix $t'_i \equiv g_{t, i}(z^G(h', S_i^{\theta'_i}, S'_{-i}))$. Since $S_i^{\theta_i}(I_i) \neq S_i^{\theta'_i}(I_i)$ and $\theta_i \cup \theta'_i \subseteq \Theta_{I_i, i}$, $I_i \in \alpha(S_i^{\theta_i}, S_i^{\theta'_i})$. Thus, OSP requires that

$$u_i(\theta_i, h, S_i^{\theta_i}, S_{-i}) \geq u_i(\theta_i, h', S_i^{\theta'_i}, S'_{-i}) \quad (36)$$

which implies

$$\theta_i + t_i \geq t'_i \quad (37)$$

and

$$u_i(\theta'_i, h, S_i^{\theta_i}, S_{-i}) \leq u_i(\theta'_i, h', S_i^{\theta'_i}, S'_{-i}) \quad (38)$$

which implies

$$\theta'_i + t_i \leq t'_i \quad (39)$$

But $\theta'_i > \theta_i$, so

$$\theta'_i + t_i > t'_i \quad (40)$$

a contradiction. This proves Lemma 1.

Lemma 2. *Suppose (G, \mathbf{S}_N) OSP-implements (f_y, f_t) and $\mathcal{P}(G, \mathbf{S}_N) = G$. Take any I_i such that $\Theta_{I_i, i}^1 \cap \Theta_{I_i, i}^0 \neq \emptyset$, and associated $a_{I_i}^*$.*

1. *If there exists $\theta_i \in \Theta_{I_i, i}^0$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then there exists t_i^0 such that:*

- (a) *For all $\theta_i \in \Theta_{I_i, i}^0$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, for all $h \in I_i$, for all S_{-i} , $g_{t, i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^0$.*

(b) For all $\theta_i \in \Theta_{I_i,i}$ such that $S_i^{\theta_i}(I_i) = a_{I_i}^*$, for all $h \in I_i$, for all S_{-i} , if $i \notin g_y(z^G(h, S_i^{\theta_i}, S_{-i}))$, then $g_{t,i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^0$.

2. If there exists $\theta_i \in \Theta_{I_i,i}^1$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then there exists t_i^1 such that:

(a) For all $\theta_i \in \Theta_{I_i,i}^1$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, for all $h \in I_i$, for all S_{-i} , $g_{t,i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^1$.

(b) For all $\theta_i \in \Theta_{I_i,i}$ such that $S_i^{\theta_i}(I_i) = a_{I_i}^*$, for all $h \in I_i$, for all S_{-i} , if $i \in g_y(z^G(h, S_i^{\theta_i}, S_{-i}))$, then $g_{t,i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^1$.

Take any type $\theta'_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta'_i}(I_i) \neq a_{I_i}^*$. Take any type $\theta''_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta''_i}(I_i) = a_{I_i}^*$. (By $\Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0 \neq \emptyset$ there exists at least one such type.) Notice that $I_i \in \alpha(S_i^{\theta'_i}, S_i^{\theta''_i})$.

By Lemma 1, $\theta'_i \notin \Theta_{I_i,i}^1$, and the game is pruned. Thus,

$$\forall h \in I_i : \forall S_{-i} : i \notin g_y(z^G(h, S_i^{\theta'_i}, S_{-i})) \quad (41)$$

Since $\theta''_i \in \Theta_{I_i,i}^0$,

$$\exists h \in I_i : \exists S_{-i} : i \notin g_y(z^G(h, S_i^{\theta''_i}, S_{-i})) \quad (42)$$

OSP requires that type θ'_i does not want to (inf-sup) deviate. Thus,

$$\begin{aligned} & \inf_{h \in I_i, S_{-i}} g_{t,i}(z^G(h, S_i^{\theta'_i}, S_{-i})) \geq \\ & \sup_{h \in I_i, S_{-i}} \{g_{t,i}(z^G(h, S_i^{\theta''_i}, S_{-i})) : i \notin g_y(z^G(h, S_i^{\theta'_i}, S_{-i}))\} \end{aligned} \quad (43)$$

OSP also requires that type θ''_i does not want to (inf-sup) deviate. This implies

$$\begin{aligned} & \inf_{h \in I_i, S_{-i}} \{g_{t,i}(z^G(h, S_i^{\theta''_i}, S_{-i})) : i \notin g_y(z^G(h, S_i^{\theta'_i}, S_{-i}))\} \\ & \geq \sup_{h \in I_i, S_{-i}} g_{t,i}(z^G(h, S_i^{\theta'_i}, S_{-i})) \end{aligned} \quad (44)$$

The RHS of Equation 43 is weakly greater than the LHS of Equation 44. The RHS of Equation 44 is weakly greater than the LHS of Equation 43. Consequently all four terms are equal. Moreover, this argument applies to every $\theta'_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta'_i}(I_i) \neq a_{I_i}^*$, and every $\theta''_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta''_i}(I_i) = a_{I_i}^*$. Since the game is pruned, θ'_i satisfies

(1b) iff $\theta_i'' \in \Theta_{I_i,i}^0$ and $S_i^{\theta_i''}(I_i) = a_{I_i}^*$. This proves part 1 of Lemma 2. Part 2 follows by symmetry; we omit the details since they involve only small changes to the above argument.

Lemma 3. *Suppose (G, \mathbf{S}_N) OSP-implements (f_y, f_t) and $\mathcal{P}(G, \mathbf{S}_N) = G$. Take any I_i such that $\Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0 \neq \emptyset$, and associated $a_{I_i}^*$. Let t_i^1 and t_i^0 be defined as before.*

1. *If there exists $\theta_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then for all $(h \in I_i, S_i, S_{-i})$, if $i \in g_y(z^G(h, S_i, S_{-i}))$, then $g_{t,i}(z^G(h, S_i, S_{-i})) \leq t_i^0 - \sup\{\theta_i \in \Theta_{I_i,i}^0 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\}$.*
2. *If there exists $\theta_i \in \Theta_{I_i,i}^1$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then for all $(h \in I_i, S_i, S_{-i})$, if $i \notin g_y(z^G(h, S_i, S_{-i}))$, then $g_{t,i}(z^G(h, S_i, S_{-i})) \leq \inf\{\theta_i \in \Theta_{I_i,i}^1 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\} + t_i^1$.*

Suppose that part 1 of Lemma 3 does not hold. Fix $(h \in I_i, S_i, S_{-i})$ such that $i \in g_y(z^G(h, S_i, S_{-i}))$ and $g_{t,i}(z^G(h, S_i, S_{-i})) > t_i^0 - \sup\{\theta_i \in \Theta_{I_i,i}^0 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\}$. Since G is pruned, we can find some $\theta_i' \in \Theta_{I_i,i}$ such that for every $\tilde{I}_i \in \{I_i' \in \mathcal{I}_i : I_i \in \psi(I_i')\}$, $S_i^{\theta_i'}(\tilde{I}_i) = S_i(\tilde{I}_i)$. Fix that θ_i' .

Fix $\theta_i'' \in \Theta_{I_i,i}^0$ such that $S_i^{\theta_i''}(I_i) \neq a_{I_i}^*$ and $\theta_i'' \geq \sup\{\theta_i \in \Theta_{I_i,i}^0 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\} - \epsilon$. Since G is pruned and $\theta_i'' \notin \Theta_{I_i,i}^1$ (by Lemma 1), it must be that $S_i^{\theta_i''}(I_i) \neq S_i^{\theta_i'}(I_i)$.

By construction, $I_i \in \alpha(S_i^{\theta_i'}, S_i^{\theta_i''})$.

OSP requires that, for all $h'' \in I_i, S_{-i}''$:

$$u_i(\theta_i'', h'', S_i^{\theta_i'}, S_{-i}'') \geq u_i(\theta_i'', h, S_i^{\theta_i'}, S_{-i}) \quad (45)$$

which entails

$$t_i^0 \geq \theta_i'' + g_{t,i}(z^G(h, S_i, S_{-i})) \quad (46)$$

which entails

$$t_i^0 - \sup\{\theta_i \in \Theta_{I_i,i}^0 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\} + \epsilon \geq g_{t,i}(z^G(h, S_i, S_{-i})) \quad (47)$$

But, by hypothesis,

$$t_i^0 - \sup\{\theta_i \in \Theta_{I_i,i}^0 : S_i^{\theta_i}(I_i) \neq a_{I_i}^*\} < g_{t,i}(z^G(h, S_i, S_{-i})) \quad (48)$$

Since this argument holds for all $\epsilon > 0$, we can pick ϵ small enough to create a contradiction. This proves part 1 of Lemma 3. Part 2 follows by symmetry.

Lemma 4. *Suppose (G, \mathbf{S}_N) OSP-implements (f_y, f_t) and $\mathcal{P}(G, \mathbf{S}_N) = G$. Take any I_i such that $|\Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0| > 1$ and associated $a_{I_i}^*$.*

1. If there exists $\theta_i \in \Theta_{I_i,i}^0$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then for all $\theta'_i \in \Theta_{I_i,i}^1$, $S_i^{\theta'_i}(I_i) = a_{I_i}^*$.
2. (Equivalently) If there exists $\theta_i \in \Theta_{I_i,i}^1$ such that $S_i^{\theta_i}(I_i) \neq a_{I_i}^*$, then for all $\theta'_i \in \Theta_{I_i,i}^0$, $S_i^{\theta'_i}(I_i) = a_{I_i}^*$.

Suppose Part 1 of Lemma 4 does not hold. Fix I_i , and choose $\theta'_i < \theta''_i$ such that $\{\theta'_i\} \cup \{\theta''_i\} \subseteq \Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0$. Fix $\theta'''_i \in \Theta_{I_i,i}^1$ such that $S_i^{\theta'''_i}(I_i) \neq a_{I_i}^*$. By Lemma 1, if $\theta'''_i \in \Theta_{I_i,i}^0$, then $S_i^{\theta'''_i}(I_i) = a_{I_i}^*$, a contradiction. Thus, $\theta'''_i \in \Theta_{I_i,i}^1 \setminus \Theta_{I_i,i}^0$, and since $\Theta_{I_i,i}^1$ dominates $\Theta_{I_i,i}^0$ in the strong set order, $\theta''_i < \theta'''_i$.

Since $\theta'_i \in \Theta_{I_i,i}^1$, there exists $h' \in I_i$ and θ'_{-i} such that $(\theta'_i, \theta'_{-i}) \in \Theta_{I_i}$ and $i \in g_y(z^G(h', S_i^{\theta'_i}, S_{-i}^{\theta'_{-i}}))$. By Lemma 2, there exists $a_i \in A_i(I_i)$ such that $a_i \neq a_{I_i}^*$ and choosing a_i ensures $i \notin y$ and $t_i = t_i^0$. Thus, by G SP

$$\theta'_i + g_{t,i}(z^G(h', S_i^{\theta'_i}, S_{-i}^{\theta'_{-i}})) \geq t_i^0 \quad (49)$$

By $\theta''_i \in \Theta_{I_i,i}^0$, there exists $h'' \in I_i$ and θ''_{-i} such that $i \notin g_y(z^G(h'', S_i^{\theta''_i}, S_{-i}^{\theta''_{-i}}))$. By Lemma 2

$$g_{t,i}(z^G(h'', S_i^{\theta''_i}, S_{-i}^{\theta''_{-i}})) = t_i^0 \quad (50)$$

By G SP, $i \in g_y(z^G(h', S_i^{\theta'''_i}, S_{-i}^{\theta'''_{-i}}))$ and $g_{t,i}(z^G(h', S_i^{\theta'''_i}, S_{-i}^{\theta'''_{-i}})) = g_{t,i}(z^G(h', S_i^{\theta'_i}, S_{-i}^{\theta'_{-i}}))$.

Notice that $I_i \in \alpha(S_i^{\theta''_i}, S_i^{\theta'''_i})$. Thus, OSP requires that θ''_i does not want to (inf-sup) deviate to θ'''_i 's strategy, which entails:

$$g_{t,i}(z^G(h'', S_i^{\theta''_i}, S_{-i}^{\theta''_{-i}})) \geq \theta''_i + g_{t,i}(z^G(h', S_i^{\theta'''_i}, S_{-i}^{\theta'''_{-i}})) \quad (51)$$

$$\begin{aligned} t_i^0 &\geq \theta''_i + g_{t,i}(z^G(h', S_i^{\theta'''_i}, S_{-i}^{\theta'''_{-i}})) \\ &> \theta'_i + g_{t,i}(z^G(h', S_i^{\theta'_i}, S_{-i}^{\theta'_{-i}})) \end{aligned} \quad (52)$$

which contradicts Equation 49.

Part 2 is the contrapositive of Part 1. This proves Lemma 4.

Lemma 5. Suppose (G, \mathbf{S}_N) OSP-implements (f_y, f_t) and $\mathcal{P}(G, \mathbf{S}_N) = G$.

For all I_i , if $|\Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0| \leq 1$ and $|A(I_i)| \geq 2$, then there exists t_i^1 and t_i^0 such that:

1. For all $\theta_i \in \Theta_{I_i,i}$, $h \in I_i$, S_{-i} :

$$(a) \text{ If } i \notin g_y(z^G(h, S_i^{\theta_i}, S_{-i})) \text{ then } g_{t,i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^0$$

(b) If $i \in g_y(z^G(h, S_i^{\theta_i}, S_{-i}))$ then $g_{t,i}(z^G(h, S_i^{\theta_i}, S_{-i})) = t_i^1$

2. If $|\Theta_{I_i,i}^1| > 0$ and $|\Theta_{I_i,i}^0| > 0$, then $t_i^1 = -\inf\{\theta_i \in \Theta_{I_i,i}^1\} + t_i^0$

By G pruned, $\Theta_{I_i,i} \neq \emptyset$.

Consider the case where $\Theta_{I_i,i}^1 = \emptyset$. Pick some $\theta'_i \in \Theta_{I_i,i}^0$ and some $h' \in I_i, S'_{-i}$. Fix $t_i^0 \equiv g_{t,i}(z^G(h', S_i^{\theta'_i}, S'_{-i}))$. Suppose there exists some $(\theta''_i, \theta''_{-i}) \in \Theta_{I_i}$ such that $f_{t,i}(\theta''_i, \theta''_{-i}) = t_i^{0''} \neq t_i^0$. Pick $h'' \in I_i$ such that $h'' \preceq z^G(\emptyset, S_i^{\theta''_i}, S_{-i}^{\theta''_{-i}})$. By Equation 10, for all $\theta_i \in \Theta_{I_i,i}$, $f_{t,i}(\theta_i, \theta''_{-i}) = t_i^{0''}$. By G pruned and $|A(I_i)| \geq 2$, we can pick $\theta_i^{0'''} \in \Theta_{I_i,i}^0$ such that $S_i^{\theta_i^{0'''}}(I_i) \neq S_i^{\theta'_i}(I_i)$. Notice that $I_i \in \alpha(S_i^{\theta_i^{0'''}}(I_i), S_i^{\theta'_i}(I_i))$. If $t_i^{0'''} > t_i^0$, then

$$u_i(\theta'_i, h', S_i^{\theta'_i}, S'_{-i}) = t_i^0 < t_i^{0''} = u_i(\theta'_i, h'', S_i^{\theta_i^{0'''}}(I_i), S_{-i}^{\theta''_{-i}}) \quad (53)$$

so $S_i^{\theta'_i}$ is not obviously dominant for (i, θ'_i) . If $t_i^{0''} < t_i^0$, then

$$u_i(\theta_i^{0'''}, h', S_i^{\theta'_i}, S'_{-i}) = t_i^0 > t_i^{0''} = u_i(\theta_i^{0'''}, h'', S_i^{\theta_i^{0'''}}(I_i), S_{-i}^{\theta''_{-i}}) \quad (54)$$

so $S_i^{\theta_i^{0'''}}$ is not obviously dominant for $(i, \theta_i^{0'''})$. By contradiction, this proves Lemma 5 for this case. A symmetric argument proves Lemma 5 for the case where $\Theta_{I_i,i}^0 = \emptyset$.

Note that, if Lemma 5 holds at some information set I_i , it holds at all information sets I'_i that follow I_i . Thus, we need only consider some earliest information set I_i^* at which $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| \leq 1$ and $|A(I_i^*)| \geq 2$.

Now we consider the case where $\Theta_{I_i^*,i}^1 \neq \emptyset$ and $\Theta_{I_i^*,i}^0 \neq \emptyset$.

At every prior information set I_i prior to I_i^* , $|\Theta_{I_i,i}^1 \cap \Theta_{I_i,i}^0| > 1$. Since $\Theta_{I_i^*,i}^1 \neq \emptyset$ and $\Theta_{I_i^*,i}^0 \neq \emptyset$, by Lemma 4, I_i^* is reached by some interval of types all taking the same action. Thus $\sup\{\theta_i \in \Theta_{I_i^*,i}^0\} = \inf\{\theta_i \in \Theta_{I_i^*,i}^1\}$.

Fix $\hat{\theta}_i \in \Theta_{I_i^*,i}^0$ such that $\hat{\theta}_i \geq \sup\{\theta_i \in \Theta_{I_i^*,i}^0\} - \epsilon$. Choose corresponding $\hat{h} \in I_i^*$ and $\hat{\theta}_{-i} \in \Theta_{I_i^*,-i}$ such that $i \notin g_y(z^G(\hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_{-i}}))$. Define $t_i^0 \equiv g_{t,i}(z^G(\hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_{-i}}))$.

Fix $\check{\theta}_i \in \Theta_{I_i^*,i}^1$ such that $\check{\theta}_i \leq \inf\{\theta_i \in \Theta_{I_i^*,i}^1\} + \epsilon$. Choose corresponding $\check{h} \in I_i^*$ and $\check{\theta}_{-i} \in \Theta_{I_i^*,-i}$ such that $i \in g_y(z^G(\check{h}, S_i^{\check{\theta}_i}, S_{-i}^{\check{\theta}_{-i}}))$. Define $t_i^1 \equiv g_{t,i}(z^G(\check{h}, S_i^{\check{\theta}_i}, S_{-i}^{\check{\theta}_{-i}}))$.

Suppose there exists some $(\theta'_i, \theta'_{-i}) \in \Theta_{I_i^*}$ such that $i \notin f_y(\theta'_i, \theta'_{-i})$ and $f_{t,i}(\theta'_i, \theta'_{-i}) = t_i^{0'} \neq t_i^0$. Since $\sup\{\theta_i \in \Theta_{I_i^*,i}^0\} = \inf\{\theta_i \in \Theta_{I_i^*,i}^1\}$, it follows that for all $\theta_{-i} \in \Theta_{I_i^*,-i}$, $\inf\{\theta_i : i \in f_y(\theta_i, \theta_{-i})\} = \inf\{\theta_i \in \Theta_{I_i^*,i}^1\}$. Thus, by Equation 10, for all $\theta_i \in \Theta_{I_i^*,i}$: $f_{t,i}(\theta_i, \theta'_{-i}) = -1_{i \in f_y(\theta_i, \theta'_{-i})} \inf\{\theta_i \in \Theta_{I_i^*,i}^1\} + t_i^{0'}$. Fix $h' \in I_i^*$ such that $h' \preceq z^G(\emptyset, S_i^{\theta'_i}, S_{-i}^{\theta'_{-i}})$.

By $A(I_i) \geq 2$, we can pick some $\theta''_i \in \Theta_{I_i^*,i}$ such that $S_i^{\theta''_i}(I_i) \neq S_i^{\hat{\theta}_i}(I_i)$. Notice that $I_i^* \in \alpha(S_i^{\theta''_i}(I_i), S_i^{\hat{\theta}_i}(I_i))$. Either $\theta''_i \in \Theta_{I_i^*}^0$ or $\theta''_i \in \Theta_{I_i^*}^1 \setminus \Theta_{I_i^*}^0$. Suppose $\theta''_i \in \Theta_{I_i^*}^0$. Suppose $t_i^{0''} > t_i^0$.

By OSP,

$$u_i(\hat{\theta}_i, \hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_i}) \geq u_i(\hat{\theta}_i, h', S_i^{\theta''_i}, S_{-i}^{\theta''_i}) \quad (55)$$

which entails

$$\begin{aligned} t_i^0 &\geq t_i^{0'} + 1_{i \in f_y(\theta''_i, \theta'_{-i})}(\hat{\theta}_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\}) \\ &\geq t_i^{0'} + 1_{i \in f_y(\theta''_i, \theta'_{-i})}(-\epsilon) \\ &\geq t_i^{0'} - \epsilon \end{aligned} \quad (56)$$

and we can pick ϵ small enough to constitute a contradiction. Suppose $t_i^{0'} < t_i^0$. By OSP

$$u_i(\theta''_i, \hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_i}) \leq u_i(\theta''_i, h', S_i^{\theta''_i}, S_{-i}^{\theta''_i}) \quad (57)$$

which entails

$$\begin{aligned} t_i^0 &\leq t_i^{0'} + 1_{i \in f_y(\theta''_i, \theta'_{-i})}(\theta''_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\}) \\ &= t_i^{0'} + 1_{i \in f_y(\theta''_i, \theta'_{-i})}(\theta''_i - \sup\{\theta_i \in \Theta_{I_i^*, i}^0\}) \leq t_i^{0'} \end{aligned} \quad (58)$$

which is a contradiction.

The case that remains is $\theta''_i \in \Theta_{I_i^*}^1 \setminus \Theta_{I_i^*}^0$. Then $i \in f_y(\theta''_i, \theta'_{-i})$ and $f_{t,i}(\theta''_i, \theta'_{-i}) = -\inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^{0'}$. Suppose $t_i^{0'} > t_i^0$. OSP requires:

$$u_i(\hat{\theta}_i, \hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_i}) \geq u_i(\hat{\theta}_i, h', S_i^{\theta''_i}, S_{-i}^{\theta''_i}) \quad (59)$$

which entails

$$\begin{aligned} t_i^0 &\geq \hat{\theta}_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^{0'} \\ &\geq t_i^{0'} - \epsilon \end{aligned} \quad (60)$$

and we can pick ϵ small enough to constitute a contradiction.

Suppose $t_i^{0'} < t_i^0$. Since $S_i^{\theta''_i}(I_i) \neq S_i^{\hat{\theta}_i}(I_i)$, either $S_i^{\check{\theta}_i}(I_i) \neq S_i^{\hat{\theta}_i}(I_i)$ or $S_i^{\check{\theta}_i}(I_i) \neq S_i^{\theta''_i}(I_i)$. Moreover, $f_{t,i}(\check{\theta}_i, \theta'_i) = -1_{i \in f_y(\check{\theta}_i, \theta'_{-i})} \inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^{0'}$.

Suppose $S_i^{\check{\theta}_i}(I_i) \neq S_i^{\hat{\theta}_i}(I_i)$. OSP requires:

$$u_i(\check{\theta}_i, h', S_i^{\check{\theta}_i}, S_{-i}^{\theta'_{-i}}) \geq u_i(\check{\theta}_i, \hat{h}, S_i^{\hat{\theta}_i}, S_{-i}^{\hat{\theta}_{-i}}) \quad (61)$$

which entails

$$\mathbf{1}_{i \in f_y(\check{\theta}_i, \theta'_{-i})}(\check{\theta}_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\}) + t_i^{0'} \geq t_i^0 \quad (62)$$

which entails

$$\mathbf{1}_{i \in f_y(\check{\theta}_i, \theta'_{-i})} \epsilon + t_i^{0'} \geq t_i^0 \quad (63)$$

and we can pick ϵ small enough to yield a contradiction. Suppose $S_i^{\check{\theta}_i}(I_i) \neq S_i^{\theta''_i}(I_i)$. By Equation 10, $f_{t,i}(\check{\theta}_i, \theta'_{-i}) = -\mathbf{1}_{i \in f_y(\check{\theta}_i, \theta'_{-i})} \inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^{0'}$, and $f_{t,i}(\theta''_i, \hat{\theta}_{-i}) = -\inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^0$. OSP requires

$$u_i(\check{\theta}_i, h', S_i^{\check{\theta}_i}, S_{-i}^{\theta'_{-i}}) \geq u_i(\check{\theta}_i, \hat{h}, S_i^{\theta''_i}, S_{-i}^{\hat{\theta}_{-i}}) \quad (64)$$

which entails

$$\mathbf{1}_{i \in f_y(\check{\theta}_i, \theta'_{-i})}(\check{\theta}_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\}) + t_i^{0'} \geq (\check{\theta}_i - \inf\{\theta_i \in \Theta_{I_i^*, i}^1\}) + t_i^0 \quad (65)$$

which entails

$$\mathbf{1}_{i \in f_y(\check{\theta}_i, \theta'_{-i})} \epsilon + t_i^{0'} \geq \epsilon + t_i^0 \quad (66)$$

which entails

$$t_i^{0'} \geq t_i^0 \quad (67)$$

a contradiction. By the above argument, for all I_i satisfying the assumptions of Lemma 5, there is a unique transfer t_i^0 for all terminal histories z passing through I_i such that $i \notin g_y(z)$. Equation 10 thus implies that there is a unique transfer t_i^1 for all terminal histories z passing through I_i such that $i \in g_y(z)$. Moreover, $t_i^1 = -\inf\{\theta_i \in \Theta_{I_i^*, i}^1\} + t_i^0$. This proves Lemma 5.

Now to bring this all together. We leave showing parts (1.d.iv) and (2.d.iv) of Definition 15 to the last. Take any $(\hat{G}, \hat{\mathbf{S}}_N)$ that OSP-implements (f_y, f_t) . Define $G \equiv \mathcal{P}(\hat{G}, \hat{\mathbf{S}}_N)$ and \mathbf{S}_N as $\hat{\mathbf{S}}_N$ restricted to G . By Proposition 2, (G, \mathbf{S}_N) OSP-implements (f_y, f_t) .

We now characterize (G, \mathbf{S}_N) . For any player i , consider any information set I_i^* such that $|A(I_i^*)| \geq 2$, and for all prior information sets $I_i' \prec I_i^*$, $|A(I_i')| = 1$. By Lemma 1,

there is a unique action $a_{I_i^*}^*$ taken by all types in $\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0$.

Either $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| > 1$ or $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| \leq 1$.

If $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| > 1$, then by Lemma 4, \tilde{G} pruned and $|A(I_i)| \geq 2$,

1. EITHER: There exists $\theta_i \in \Theta_{I_i^*,i}^0$ such that $S_i^{\theta_i}(I_i^*) \neq a_{I_i^*}^*$, and for all $\theta'_i \in \Theta_{I_i^*,i}^1$, $S_i^{\theta'_i}(I_i^*) = a_{I_i^*}^*$.
2. OR: There exists $\theta_i \in \Theta_{I_i^*,i}^1$ such that $S_i^{\theta_i}(I_i^*) \neq a_{I_i^*}^*$, and for all $\theta'_i \in \Theta_{I_i^*,i}^0$, $S_i^{\theta'_i}(I_i^*) = a_{I_i^*}^*$.

In the first case, then by Lemma 2, there is some \bar{t}_i such that, for all (S_i, S_{-i}) , for all $h \in I_i$, if $i \notin g_y(z^G(h, S_i, S_{-i}))$, then $g_{t,i}(z^G(h, S_i, S_{-i})) = \bar{t}_i$. Moreover, we can define a ‘going transfer’ at all information sets I'_i such that $I_i^* \preceq I'_i$:

$$\tilde{t}_i(I'_i) \equiv \min_{I''_i: I_i^* \preceq I''_i \preceq I'_i} [\bar{t}_i - \sup\{\theta_i \in \Theta_{I''_i,i}^0 : S_i^{\theta_i}(I''_i) \neq a_{I''_i}^*\}] \quad (68)$$

Notice that this function falls monotonically as we move along the game tree; for any I'_i, I''_i such that $I'_i \preceq I''_i$, $\tilde{t}_i(I'_i) \geq \tilde{t}_i(I''_i)$. Moreover, by construction, at any I'_i, I''_i such that $I'_i \prec I''_i$, $\tilde{t}_i(I'_i) > \tilde{t}_i(I''_i)$, and there does not exist I'''_i such that $I'_i \prec I'''_i \prec I''_i$, then there exists $a \in A(I''_i)$ that yields $i \notin y$, and by Lemma 2 this yields transfer \bar{t}_i . We define A^q to include all such quitting actions; i.e. A^q is the set of all actions such that:

1. $a \in I_i$ for some $I_i \in \mathcal{I}_i$
2. For all z such that $a \in \psi_i(z)$: $i \notin g_y(z)$ and $g_{t,i}(z) = \bar{t}_i$

Lemma 3 and SP together imply that, at any terminal history z , if $i \in g_y(z)$, then

$$g_{t,i}(z) = \inf_{I_i: I_i^* \preceq I_i \prec z} \tilde{t}_i(I_i) \quad (69)$$

This holds because, if $g_{t,i}(z) < \inf_{I_i: I_i^* \preceq I_i \prec z} \tilde{t}_i(I_i)$, then type θ_i such that $\bar{t}_i - \inf_{I_i: I_i^* \preceq I_i \prec z} \tilde{t}_i(I_i) < \theta_i < \bar{t}_i - g_{t,i}(z)$ could profitably deviate to play $a \in A^0$ at information set I_i^* .

In the second case, then by Lemma 2, there is some \bar{t}_i such that, for all (S_i, S_{-i}) , for all $h \in I_i^*$, if $i \in g_y(z^G(h, S_i, S_{-i}))$, then $g_{t,i}(z^G(h, S_i, S_{-i})) = \bar{t}_i$. Moreover, we can define a ‘going transfer’ at all information sets I'_i such that $I_i^* \preceq I'_i$:

$$\tilde{t}_i(I'_i) \equiv \min_{I''_i: I_i^* \preceq I''_i \preceq I'_i} [\bar{t}_i + \inf\{\theta_i \in \Theta_{I''_i,i}^0 : S_i^{\theta_i}(I''_i) \neq a_{I''_i}^*\}] \quad (70)$$

We define A^1 symmetrically for this second case.

Part (1.d.iii) and (and its analog in Clause 2) of Definition 15 follow from Lemma 4. The above constructions suffice to prove Theorem 3 for cases where $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| > 1$. Cases where $|\Theta_{I_i^*,i}^1 \cap \Theta_{I_i^*,i}^0| \leq 1$ are dealt with by Lemma 5.

Now for the last piece: We prove that parts (1.d.iv) and (and its analog in Clause 2) of Definition 15 hold. The proof of part (1.d.iv) is as follows: Suppose we are facing Clause 1 of Definition 15, and for some I_i , $|A(I_i') \setminus A^0| > 1$. By part (1.d.iii), we know that the going transfer \tilde{t}_i can fall no further. Since G is pruned and $|A(I_i') \setminus A^0| > 1$, there exist two distinct types of i , $\theta_i, \theta'_i \in \Theta_{I_i',i}$, who do not quit at I_i' , and take different actions. Since neither quits at I_i' and the going transfer falls no further, there exist $\theta_{-i}, \theta'_{-i} \in \Theta_{I_i',-i}$ such that $i \in f_y(\theta_i, \theta_{-i})$ and $i \in f_y(\theta'_i, \theta'_{-i})$. So there exist $(h \in I_i', S_{-i})$ and $(h' \in I_i', S'_{-i})$ such that

$$i \in g_y(h, S_i^{\theta_i}, S_{-i}) \quad (71)$$

$$i \in g_y(h', S_i^{\theta'_i}, S'_{-i}) \quad (72)$$

$$g_{t,i}(h, S_i^{\theta_i}, S_{-i}) = g_{t,i}(h', S_i^{\theta'_i}, S'_{-i}) = \tilde{t}_i(I_i') \quad (73)$$

WLOG suppose $\theta_i < \theta'_i$. Suppose that there does not exist $a \in A(I_i')$ such that, for all z such that $a \in \psi_i(z)$, $i \in g_y(z)$. Then there must exist $(h'' \in I_i', S''_{-i})$ such that

$$i \notin g_y(h'', S_i^{\theta'_i}, S''_{-i}) \quad (74)$$

$$g_{t,i}(h'', S_i^{\theta'_i}, S''_{-i}) = \bar{t}_i \quad (75)$$

Note that $I_i' \in \alpha(S_i^{\theta_i}, S_i^{\theta'_i})$. But then $S_i^{\theta'_i}$ is not obviously dominant, a contradiction, since

$$\begin{aligned} u_i^{\tilde{G}}(h'', S_i^{\theta'_i}, S''_{-i}, \theta'_i) &= \bar{t}_i \leq \theta_i + \tilde{t}_i(I_i') \\ &< \theta'_i + \tilde{t}_i(I_i') = u_i^{\tilde{G}}(h, S_i^{\theta_i}, S_{-i}, \theta_i) \end{aligned} \quad (76)$$

(The first inequality holds because of type θ_i 's incentive constraint.) This shows that part (1.d.iv) of Definition 15 holds. Its analog in Clause 2 is proved symmetrically. \square

Proposition 7. *If G is a personal-clock auction, then there exist \mathbf{S}_N and f_y such that*

(G, \mathbf{S}_N) OSP-implements f_y .

Proof. Take any personal-clock auction G . For any i and any θ_i , the following strategy S_i is obviously dominant:

1. If i encounters an information set consistent with Clause 1 of Definition 15, then, from that point forward:
 - (a) If $\theta_i + \tilde{t}_i(I_i) > \bar{t}_i$ and there exists $a \in A(I_i) \setminus A^q$, play $a \in A(I_i) \setminus A^q$.
 - i. If $|A(I_i) \setminus A^q| > 1$, then play $a \in A(I'_i)$ such that: For all z such that $a \in \psi_i(z)$: $i \in g_y(z)$.
 - (b) Else play some $a \in A^q$.⁶¹
2. If i encounters an information set consistent with Clause 2 of Definition 15, then, from that point forward:
 - (a) If $\theta_i + \bar{t}_i < \tilde{t}_i(I_i)$ and there exists $a \in A(I_i) \setminus A^q$, play $a \in A(I_i) \setminus A^q$.
 - i. If $|A(I_i) \setminus A^q| > 1$, then play $a \in A(I'_i)$ such that: For all z such that $a \in \psi_i(z)$: $i \notin g_y(z)$.
 - (b) Else play some $a \in A^q$.

The above strategy is well-defined for any agent in any personal-clock auction, by inspection of Definition 15.

Consider any deviating strategy S'_i . At any earliest point of departure, the agent will have encountered an information set consistent with either Clause 1 or Clause 2 of Definition 15. Suppose that the agent has encountered an information set covered by Clause 1.

Take some earliest point of departure $I_i \in \alpha(S_i, S'_i)$. Notice that, by (1.a) of Definition 15, *no matter what strategy i plays*, conditional on reaching I_i , either agent i is not in the allocation and receives \bar{t}_i , or agent i is in the allocation and receives a transfer $\hat{t}_i \leq \tilde{t}_i(I_i)$.

Suppose $\theta_i + \tilde{t}_i(I_i) > \bar{t}_i$. Note that under S_i , conditional on reaching I_i , the agent either is not in the allocation and receives \bar{t}_i , or is in the allocation and receives a transfer strictly above $\bar{t}_i - \theta_i$. If $S'_i(I_i) \in A_q$ (*i.e.* if agent i quits), then the best outcome under S'_i is no better than the worst outcome under S_i . If $S'_i(I_i) \notin A_q$, then, since $S'_i(I_i) \neq S_i(I_i)$, $|A(I_i) \setminus A^q| > 1$. Then, by (1.d.iii) of Definition 15, \tilde{t}_i will fall no further. So $S_i(I_i)$ guarantees that i will be in the allocation and receive transfer $\tilde{t}_i(I_i)$. But, by (1.a) of

⁶¹If $|A^q \cap A(I_i)| > 1$, the agent chooses deterministically but arbitrarily.

Definition 15, the best possible outcome under S'_i conditional on reaching I_i is no better, so the obvious dominance inequality holds.

Suppose $\theta_i + \tilde{t}_i(I_i) \leq \bar{t}_i$. Then, under S_i , conditional on reaching I_i , agent i is not in the allocation and has transfer \bar{t}_i .⁶² However, under S'_i , either the outcome is the same, or agent i is in the allocation for some transfer $\hat{t}_i \leq \tilde{t}_i(I_i) \leq \bar{t}_i - \theta_i$. Thus, the best possible outcome under S'_i is no better than the worst possible outcome under S_i , and the obvious dominance inequality holds.

The argument proceeds symmetrically for Clause 2.

Notice that the above strategies result in *some* allocation and *some* payments, as a function of the type profile. We define these to be (f_y, f_t) , such that G OSP-implements (f_y, f_t) . □

C Alternative Empirical Specifications (for online publication)

Here we report alternative empirical specifications for the experiment.

A natural measure of errors would be to take the sum, for $k = 1, 2, 3, 4$, of the absolute difference between the k th highest bid and the k th highest value. However we do not observe the highest bid under **AC**, and we often do not observe the highest bid under **AC+X**. We could instead take the sum for $k = 2, 3, 4$ of the absolute difference between the k th highest bid and the k th highest value, averaged as before in five-round blocks. Table 6 reports the results.

Another measure of errors would be to take the sum of the absolute difference between each bidder's bid and *that bidder's* value, dropping all highest bidders for symmetry. Table 7 reports the results.

Table 8 reports the results of Table 4, except that the p -values are calculated using the Wilcoxon rank-sum test.

29.0% of rank-order lists are incorrect under **SP-RSD**. 2.6% of choices are incorrect under **OSP-RSD**. However, this is not a fair comparison; rank-order lists mechanically allow us to spot more errors than single choices. To compare like with like, we compute the proportion of incorrect choices we would have observed, if subjects played **OSP-RSD** as though they were implementing the submitted rank-order lists for **SP-RSD**. This is a cautious measure; it counts errors under **SP-RSD** only if they alter the outcome. Table

⁶²By (1.d.ii) of Definition 15, either i will have quit in the past, or will have an opportunity to quit now, which he exercises.

Table 6: mean(sum(abs(k th bid - k th value))), for $k = 2, 3, 4$

Format	Rounds	SP	OSP	p -value
Auction	1-5	32.63 (4.64)	9.89 (1.89)	< .001
	6-10	16.28 (2.73)	5.53 (0.91)	.001
+X Auction	1-5	17.04 (3.70)	6.18 (1.06)	.011
	6-10	14.21 (3.70)	4.74 (0.75)	.022

For each auction, we sum the absolute differences between the k th bid and the k th value, for $k = 2, 3, 4$. We then take the mean of this over each 5-round block. We then compute standard errors counting each group's 5-round mean as a single observation. (18 observations per cell.) p -values are computed using a two-sample t -test, allowing for unequal variances.

Table 7: mean(sum(abs(i 's bid - i 's value))), dropping highest bidders

Format	Rounds	SP	OSP	p -value
Auction	1-5	35.13 (5.20)	10.18 (1.88)	< .001
	6-10	15.46 (2.85)	4.89 (0.72)	.002
+X Auction	1-5	17.88 (4.10)	5.58 (1.01)	.009
	6-10	14.20 (3.72)	4.64 (0.83)	.022

For each auction, we sum the absolute differences between each bidder's bid and their value, dropping the highest bidder. We then take the mean of this over each 5-round block. We then compute standard errors counting each group's 5-round mean as a single observation. (18 observations per cell.) p -values are computed using a two-sample t -test, allowing for unequal variances.

Table 8: Proportion of serial dictatorships not ending in dominant strategy outcome, p -values calculated using Wilcoxon rank-sum test

	SP	OSP	p -value
Rounds 1-5	43.3%	7.8%	.0001
Rounds 6-10	28.9%	6.7%	.0010

This is the same as Table 4, except that the p -values are calculated using the Wilcoxon rank-sum test.

9 reports the results.

Table 9: Proportion of incorrect choices under serial dictatorship: SP (imputed) vs OSP (actual)

	SP	OSP	p -value
Rounds 1-5	17.8% (3.5%)	2.6% (1.1%)	< .001
Rounds 6-10	10.7% (2.0%)	2.6% (1.3%)	.002
p -value	.078	1.000	

For each group in **SP-RSD**, for each period, we simulate the three choices that we would have observed under **OSP-RSD**. For each group, for each 5-round block, we record the proportion of choices that are incorrect. We then compute standard errors counting each group-block pair as a single observation. (18 observations per cell.) When comparing SP to OSP, we compute p -values using a two-sample t -test. When comparing early to late rounds of the same game, we compute p -values using a paired t -test. In the sample for **OSP-RSD**, there are 7 incorrect choices in the first five rounds and 7 incorrect choices in the last five rounds.

In Table 10, we compute the mean difference between the second-highest bid and the second-highest value, by auction format and by five-round blocks. This summarizes the average *direction* of deviations (positive or negative) from equilibrium play. I had no prior hypothesis about these outcome variables, but report this analysis for completeness.

Table 11 breaks down the rate that subjects submit incorrect rank-order lists in one-shot **SP-RSD** by demographic category.

Table 10: mean(2nd bid - 2nd value)

Format	Rounds	SP	OSP
Auction	1-5	-1.59 (2.07)	0.58 (0.77)
	6-10	4.22*** (1.28)	0.78* (0.40)
+X Auction	1-5	1.62** (0.70)	0.15 (0.47)
	6-10	2.97*** (0.97)	0.02 (0.31)

For each group, we take the mean difference between the second-highest bid and the second-highest value over each 5-round block. We then compute standard errors counting each group's 5-round mean as a single observation. (18 observations per cell, standard errors in parentheses.) We then run a two-sided t -test for the null that the value in each cell has zero mean. * denotes $p < .10$, ** denotes $p < .05$, and *** denotes $p < .01$.

Table 11: Incorrect lists in one-shot **SP-RSD**, by demographic categories

	Yes	No	<i>p</i> -value
Economics major	46%	31%	.131
STEM major	32%	34%	.876
Took a course in game theory	34%	33%	.876
Male	37%	28%	.214

This table displays the proportion of subjects who submitted incorrect rank-order lists in one-shot **SP-RSD**, by self-reported demographic categories. There are 192 subjects in this treatment. *p*-values are computed with Fisher’s exact test.

D Quantal Response Equilibrium (for online publication)

Quantal response equilibrium (QRE) is defined for normal form games (McKelvey and Palfrey, 1995). Agent quantal response equilibrium (AQRE) adapts QRE to extensive form games (McKelvey and Palfrey, 1998); agents play a perturbed best response at each information set, given correct beliefs about the (perturbed) play of other agents and their future selves.

Standard specifications of AQRE predict extreme under-bidding in ascending auctions with fine bid increments. I here provide a simple proof that, for any private-value ascending auction, there exist bid increments fine enough that logit-AQRE predicts that all bids will be close to 0 with probability close to 1. The proof also works for a class of related models, that allow agents to have arbitrary beliefs about the strategies of other agents and their future selves.

Definition 17. *An equilibrium with τ -logit errors specifies, for every information set, some (arbitrary) beliefs about opponent strategies and the agent’s own continuation strategy. Given those beliefs, let $\{v_1, \dots, v_J\}$ denote the implied continuation values for actions $\{a_1, \dots, a_J\}$ at information set I_i . The probability that agent i chooses a_j at I_i is equal to $\frac{e^{\tau v_j}}{\sum_{l=1}^J e^{\tau v_l}}$ for some $\tau \geq 0$, where this randomization is independent across information sets and agents.*

This contains logit-AQRE (McKelvey and Palfrey, 1998) as a special case. It also includes naïve logit models, such that agents believe that their opponents and their future selves play dominant strategies.

We now consider the following environment: There is a single object and some finite set of bidders, with types (private values) in the interval $[0, 1]$ and quasi-linear utility. Types are drawn according to some joint cdf $F : [0, 1]^{|N|} \rightarrow [0, 1]$.

Definition 18. *In a k -increment ascending auction, the price rises step by step*

through the set $\{\frac{0}{k}, \frac{1}{k}, \frac{2}{k}, \dots\}$. Agents observe only the clock price, and whether the auction has ended. At each price, active agents (in some arbitrary order) choose In or Out. Agents who choose Out become inactive; they make no payments and do not receive the object. At any point when only one agent is active, the auction ends, and the active agent wins the object at the current clock price.

Proposition 8. For any $\epsilon > 0$ and any τ , there exists \bar{k} such that, for all $k \geq \bar{k}$, in any equilibrium with τ -logit errors of the k -increment ascending auction, the probability that the closing price is above ϵ is less than ϵ .

Proof. Let $M_k^\epsilon := \max\{j \in \mathbb{Z} \text{ such that } \frac{j}{k} \leq \epsilon\}$; this is the index of the highest bid increment below ϵ . Take any agent i , any type θ_i , any k -increment auction, and any equilibrium with τ -logit errors. Let $v_{i,j,k,\theta_i}^{\text{In}}$ denote the continuation value (given i 's arbitrary beliefs) of playing In at clock price $\frac{j}{k}$.

$$\begin{aligned} & P([i \text{ does not quit before } \epsilon] \cap [\text{some } j \in N \setminus i \text{ does not quit before } \epsilon]) \\ & \leq P(i \text{ does not quit before } \epsilon \mid \text{some } j \in N \setminus i \text{ does not quit before } \epsilon) \\ & = \prod_{j=0}^{M_k^\epsilon} \frac{e^{\tau v_{i,j,k,\theta_i}^{\text{In}}}}{e^{\tau v_{i,j,k,\theta_i}^{\text{In}}} + e^0} \leq \prod_{j=0}^{M_k^\epsilon} \frac{e^\tau}{e^\tau + e^0} \quad (77) \end{aligned}$$

The last inequality holds because, for *any* beliefs, i 's payoff from the k -increment ascending auction is bounded above by 1. Moreover,

$$\begin{aligned} & P(\text{closing price above } \epsilon) \\ & = P\left(\bigcup_{i \in N} [[i \text{ does not quit before } \epsilon] \cap [\text{some } j \in N \setminus i \text{ does not quit before } \epsilon]]\right) \\ & \leq \sum_{i \in N} P([i \text{ does not quit before } \epsilon] \cap [\text{some } j \in N \setminus i \text{ does not quit before } \epsilon]) \\ & \leq |N| \prod_{j=0}^{M_k^\epsilon} \frac{e^\tau}{e^\tau + e^0} \quad (78) \end{aligned}$$

$\lim_{k \rightarrow \infty} M_k^\epsilon = \infty$, so $\lim_{k \rightarrow \infty} \prod_{j=0}^{M_k^\epsilon} \frac{e^\tau}{e^\tau + e^0} = 0$. Consequently, we can pick \bar{k} large enough that, for all $k \geq \bar{k}$, the right-hand side of the above equation is less than ϵ , which completes the proof. \square

Table 12: AQRE computed for simplified auctions

τ	mean(abs(2nd bid - 2nd value))	
	2nd-Price	Ascending
0	4.57	11.61
2^{-4}	4.34	11.52
2^{-2}	3.79	11.06
1	2.69	9.32
2^2	1.63	5.91
2^4	0.91	2.67

As Proposition 8 indicates, in any ascending auction with fine enough bid increments, models such as logit-AQRE or naïve logit predict extreme underbidding.

In the auctions in my experiment, the range of types is \$120 and bid increments are 25 cents. It is not feasible to compute AQRE for these auctions. Players have 481 possible types and there are 601 possible bids. Consequently, the pure strategy space in **2P** and **2P+X** is of the order of $601^{481} \approx 10^{1336}$, and larger yet in **AC** and **AC+X**, since there are multiple payoff-equivalent strategies of each kind (specifying different actions at information sets that are never reached).

However, I compute logit-AQRE for the following simpler games: There are 4 agents, with types drawn independently and uniformly at random from 0.25 to 20 (in 0.25 increments). Agents either play a second-price auction with 80 possible bids (equal to the types), or an ascending clock auction with 80 clock prices.

I compute symmetric logit-AQRE for $\tau \in \{0, 2^{-4}, 2^{-2}, 1, 2^2, 2^4\}$, and sample 10^6 plays for each game-parameter pair. Table 12 displays the average *absolute* difference between the second-highest bid and the second-highest value, by τ and auction format. For every computed parameter value, AQRE predicts larger errors in the ascending auction than in the second-price auction, which is the opposite of the effect in the experimental data.

For the simplified auctions that I computed, Table 13 displays the average difference between the second-highest bid and the second-highest value predicted by AQRE. Comparing Table 12 and Table 13 indicates that the large errors in ascending auctions under AQRE are driven almost entirely by under-bidding.

As Figure 2a and Figure 2b illustrate, there is no evidence of systematic under-bidding in **AC** and **AC+X**. Table 10 shows this rigorously.

For further robustness, I compute a ‘naïve’ alternative to AQRE: Agents make logistic errors at each information set, while believing that their opponents (and their future selves) play their dominant strategies⁶³. Table 14 displays the average *absolute* difference

⁶³In ascending auctions, there are multiple payoff-equivalent dominant strategies, which prescribe

Table 13: Directional Errors under AQRE

τ	mean(2nd bid - 2nd value)	
	2nd-Price	Ascending
0	0.01	-11.61
2^{-4}	-0.33	-11.52
2^{-2}	-0.83	-11.06
1	-0.75	-9.32
2^2	-0.32	-5.91
2^4	-0.06	-2.65

Table 14: Naïve logistic errors computed for simplified auctions

τ	mean(abs(2nd bid - 2nd value))	
	2nd-Price	Ascending
0	4.57	11.61
2^{-4}	4.35	11.59
2^{-2}	3.89	11.52
1	2.86	10.71
2^2	1.68	6.02
2^4	0.91	2.43

between the second-highest bid and the second-highest value, by τ and auction format. This model also predicts larger mistakes in ascending auctions than in second-price auctions.

The preceding results suggest that models in which agents make payoff-sensitive mistakes *independently* at every information set are substantially at odds with the data. They predict that ascending auctions with fine increments will induce extreme underbidding, and that errors should be larger in ascending auctions than in second-price auctions. One could invent new theories that permit non-independent errors across information sets or different accuracy parameters for different games, but this introduces many new degrees of freedom, and may have little testable content.

different actions at prices above the agents' values. For the purpose of these computations, I assume that agents believe that they will quit at all such information sets.