

A revelation principle for obviously strategy-proof implementation

Citation for published version (APA):

Mackenzie, A. (2018). *A revelation principle for obviously strategy-proof implementation*. Maastricht University, Graduate School of Business and Economics. GSBE Research Memoranda No. 014
<https://doi.org/10.26481/umagsb.2018014>

Document status and date:

Published: 08/05/2018

DOI:

[10.26481/umagsb.2018014](https://doi.org/10.26481/umagsb.2018014)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Andrew Mackenzie

**A revelation principle for
obviously strategy-proof
implementation**

RM/18/014

GSBE

Maastricht University School of Business and Economics
Graduate School of Business and Economics

P.O. Box 616
NL- 6200 MD Maastricht
The Netherlands

A REVELATION PRINCIPLE FOR OBVIOUSLY STRATEGY-PROOF IMPLEMENTATION

ANDREW MACKENZIE*

This draft: April 9, 2018

Abstract

We prove that if a stochastic (social choice) rule has an *obviously strategy-proof* (OSP) implementation (Li, 2016), then it has such an implementation through a *randomized round table mechanism*, where the administrator randomly selects a game form in which the agents take turns making public announcements about their private information. When restricted to deterministic rules, our result improves upon other recent revelation principles by relaxing all recall requirements and by allowing all game trees compatible with normal forms (Alós-Ferrer and Ritzberger, 2016); we also establish robustness to player randomization using novel solution concepts involving mixed strategies and behavioral strategies. We use our result to provide a justification for ordinal mechanisms in the spirit of Carroll (2017), and we provide a simple characterization of the deterministic rules with OSP-implementations using deterministic round table mechanisms and ordinary *strategy-proofness*.

Keywords: obvious strategy-proofness, revelation principle, randomized round table mechanism

1 Introduction

1.1 Overview

A group of agents who wish to condition a decision on their collective information according to some *rule* may face a problem familiar in politics and economics: this collective information is known to nobody in its entirety. Such a situation can merit the design of an institution called a *game form* (or *mechanism*), which specifies an unambiguous procedure by which the agents make choices to ultimately determine their fate. The plausibility of such an institution succeeding is precisely the plausibility of the agents using what they individually know to make choices that lead to, or *implement*, the desired outcome, and this is typically articulated using a *solution concept*: a list of plausible equilibria for every game. In response to experimental evidence challenging the universal

*Department of Economics, Maastricht University, Maastricht, the Netherlands. Email: a.mackenzie@maastrichtuniversity.nl. I thank Gabriel Carroll, Yannai Gonczarowski, Hervé Moulin, Marek Pycia, Yangwei Song, and William Thomson for their comments. I am particularly grateful to Shengwu Li for his thorough feedback and suggestions.

plausibility of *dominant strategy equilibrium*, the more demanding notion of *obviously dominant strategy equilibrium* was recently proposed (Li, 2016), and the objective of this article is to establish a *revelation principle* promising that any investigation of implementation with this new solution concept—what is called *obviously strategy-proof (OSP)* implementation—can be safely restricted to a class of canonical game forms.

To understand the new demands of **OSP**-implementation, recall that one strategy dominates another if it yields an outcome that is at least as desirable no matter what happens. If there are many things that might happen, then, calculating that one strategy dominates another is a rather taxing exercise. On the other hand, if one compares the two strategies only when confronted with the choice of continuing with one or the other, and at this point *every* outcome the first strategy might yield is at least as desirable as *every* outcome the second might, then a thorough case-by-case calculation is unnecessary; the first strategy *obviously dominates* the second. In this way, a strategy might be dominant but not obviously so, and an equilibrium might consist of strategies that are dominant but not obviously so.

Due to these new demands, game forms that were adequate for dominant strategy implementation may fall short for **OSP**-implementation. Indeed, though the classic revelation principle promises that any investigation of implementation in dominant strategies can be safely restricted to *direct mechanisms*, where agents simultaneously report what they know (Gibbard, 1973; Myerson, 1981), this is not true for **OSP**-implementation (Li, 2016). This point is illustrated by a variety of examples:

- for the allocation of one object and money, the *second-price auction* (Vickrey, 1961) can be **OSP**-implemented (through an ascending clock auction), but not through its associated direct mechanism (Li, 2016);
- for object allocation, any *dual dictatorship rule*¹ can be **OSP**-implemented, but not through its associated direct mechanism (Trojan, 2016);
- for random object allocation, *randomized serial dictatorship* can be “random **OSP**-implemented,” but not through its associated direct mechanism (Pycia and Trojan, 2017);
- for the marriage problem, if preferences are “very homogeneous” for one side of the market, then *deferred acceptance* (Gale and Shapley, 1962) can be **OSP**-implemented, but not through its associated direct mechanism (Ashlagi and Gonczarowski, 2016); and
- for social choice among ordered alternatives when preferences are single-peaked, any *dictatorship with safeguards against extremism rule* has an **OSP**-implementation, but not through its associated direct mechanism (Bade and Gonczarowski, 2017; see also Arribillaga, Massó, and Neme, 2017).

Thus an analyst may like to know: is there *some* canonical class of game forms to which a mechanism designer interested in **OSP**-implementation *can* safely restrict attention? The purpose of this paper is to provide the strongest possible reassurance that this is indeed the case, and to do so at a level of generality that covers the above applications and many others.

¹In a *dual dictatorship rule*, property rights are defined using a priority structure in a narrow class, then *Gale’s top trading cycles* (described in Shapley and Scarf, 1974) is applied.

The usefulness of such a revelation principle to the analysis of **OSP**-implementation has in fact already been demonstrated. Independently of this paper, other recent work has developed several versions of the following result and applied it in various settings: if a deterministic rule can be **OSP**-implemented in pure strategies through a game form satisfying some notion of recall, then it can be **OSP**-implemented in pure strategies through a game form where agents take turns publicly announcing some of what they do and do not know (Ashlagi and Gonczarowski, 2016; Pycia and Troyan, 2017; Bade and Gonczarowski, 2017). In such a game form, an agent may not reveal everything he knows at once, but as the agents confide in one another turn after turn, group trust is slowly built in the sense that it eventually becomes apparent to all that it is safe to share the truth. We refer to these game forms as *round table mechanisms*. See Section 1.3 for a detailed comparison of our Proposition 2 to these other results.

Unfortunately, the above result is of limited use in the analysis of *stochastic rules*, where a group would like to condition a probabilistic distribution of outcomes on their collective information—for two reasons. First, the result only applies if the analyst restricts attention to game forms where the players’ choices determine lotteries, and this modeling approach faces a conceptual issue: if a terminal history assigned a lottery is appropriately expanded into a history where the administrator randomly selects a terminal history assigned a sure outcome, then a strategy that was obviously dominant may no longer be. In other words, when a player contemplates what may happen when comparing strategies, he must consider randomization during play in a different manner from randomization after play. Second, even if this conceptual issue is set aside, the modeling approach requires all desired randomization to occur after play concludes instead of over the course of play, which is a serious restriction for **OSP**-implementation. This restriction is sharply illustrated with random serial dictatorship for the assignment of indivisible objects to agents: while this rule has a “random **OSP**-implementation” where the administrator randomizes at the start of play (Pycia and Troyan, 2017), it has no **OSP**-implementation where the desired lotteries are determined by play (Theorem 2). When investigating the **OSP**-implementation of stochastic rules, then, it is *not* safe to restrict attention to round table mechanisms.

Fortunately, there *is* in fact a canonical class to which the analyst can safely restrict attention for a formal notion of random **OSP**-implementation. In particular, our main result states that for the analysis of stochastic rules, it is safe to restrict attention to what we call *randomized round table mechanisms*, where the administrator acts once at the start of play, randomly selecting a round table mechanism (Theorem 1). Familiar examples of randomization at the start of a mechanism include the selection of a priority order for random serial dictatorship and the selection of school priorities for students in school choice. We emphasize that this revelation principle still requires the analyst to consider the entire class of randomized round table mechanisms, in contrast to the classic revelation principle for dominant strategy implementation, which allows the analyst to focus on a single canonical game form; this point is also raised by Arribillaga, Massó, and Neme (2017).

The precise statement of our main result is particularly general; when applied to deterministic rules, it improves on the other revelation principles in several ways. Most notably, it has thus far not been known whether or not there are any rules that can only be **OSP**-implemented through game forms that *violate* recall requirements. In situations where players are individuals, such a game form might require someone to forget a previous action—or even to be unsure if he has already been called to play—and it perhaps

goes without saying that the institutionalized altering of the contents of an individual’s memory raises serious ethical (and practical!) concerns. On the other hand, in situations where players are groups (such as organizations, firms, or households), such a rule could be **OSP**-implemented if various representatives of a group are sometimes called to play without knowledge of whether or not other representatives have been called and how they have played—say, if the mechanism’s administrator decides who to interview and what to ask based on previous interview responses (Isbell, 1957). Because we drop all information set restrictions, our result implies that there are no such rules: there is no need for technology that alters memories, and it is always sufficient to summon a single representative from each group to play.²

1.2 Further results

After establishing our revelation principle and the necessity of this generalization, we provide three further results related to **OSP**-implementation and canonical classes of mechanisms.

First, we generalize the model to accommodate randomization by the players in addition to the administrator. At first glance, it is ambiguous whether this would expand or shrink the collection of rules which can be **OSP**-implemented. On the one hand, perhaps more rules can be implemented simply because there are additional strategies that players have access to. On the other hand, perhaps fewer rules can be implemented: when an agent contemplates what might happen when comparing strategies, he has more to worry about when he and his opponents might randomize, reducing occurrences of obvious dominance. This latter point is especially a concern in game forms that violate recall assumptions: under the interpretation that representatives of a group are called to play, behavioral strategies allow a group to realize choices that, without randomization, would require coordination between representatives that is forbidden by the structure of the game form.

Unfortunately, our proof approach is not able to accommodate probabilistic notions of **OSP**-implementation, even if it is almost-sure (Example 1). That said, using a sure notion of **OSP**-implementation, we establish that if a deterministic rule has an **OSP**-implementation in any kind of strategies (pure, mixed, or behavioral), then it has an **OSP**-implementation in pure strategies through a round table mechanism—even if a group entertains the possibility that it, or its opponents, may later discover a way to coordinate choices in a manner that is forbidden by the laws of the mechanism (Theorem 3). Altogether, while player randomization does not help, the contemplation of player randomization does not hurt; in this way our reassurance that it is safe to restrict attention to round table mechanisms is particularly robust.

We then use our revelation principle to prove a result in the spirit of Carroll (2017) for settings where there is a finite collection of outcomes, and where agents have von Neumann-Morgenstern types compatible with strict rankings of outcomes. In particular, we prove that if a stochastic rule has an **OSP**-implementation, then it is ordinal, and moreover has a round table implementation where agents make public announcements simply about their rankings of deterministic *outcomes* (Theorem 4). Whereas the justi-

²Our main result also improves on previous revelation principles by using the largest class of game trees for which choices determine outcomes (Alós-Ferrer and Ritzberger, 2016); as a side contribution, we provide a new description of this class (Proposition 1).

fication for ordinal mechanisms provided by [Carroll \(2017\)](#) relies on interdependence of agents’ preferences over lotteries, ours does not due to our stronger solution concept.

Finally, we consider ordinary *strategy-proofness* (**SP**) in the context of round table mechanisms. As an immediate corollary to [Theorem 2](#), if no round table mechanism provides an **SP**-implementation for a deterministic rule, then the rule has no **OSP**-implementation at all. Remarkably, the converse is true as well, providing a characterization of the deterministic rules with **OSP**-implementations: a deterministic rule has an **OSP**-implementation if and only if it has a **SP**-implementation through a round table mechanism ([Theorem 5](#)).³ Thus any investigation of **OSP**-implementation for a particular deterministic rule can be performed entirely using a familiar solution concept and a simple class of game forms.

1.3 Related revelation principles

In this section, we provide a detailed comparison of our [Proposition 2](#) for deterministic rules to the other revelation principles in the literature.

A revelation principle takes the following form: if a rule can be implemented through a game form in class Γ_{BIG} , then it can be implemented through one in class $\Gamma_{SMALL} \subseteq \Gamma_{BIG}$. Such a statement is made logically stronger when Γ_{BIG} is made larger, as well as when Γ_{SMALL} is made smaller. Independently of this work, a series of closely-related revelation principles for obvious strategy-proofness were established in recent working papers ([Ashlagi and Gonczarowski, 2016](#); [Pycia and Troyan, 2017](#); [Bade and Gonczarowski, 2017](#)), and this section is dedicated to discussing the relationship between our revelation principle and these three.⁴ We remark that in each of these other papers, the revelation principle is not the primary contribution, but rather plays the role of a lemma toward the main result.

[Ashlagi and Gonczarowski \(2016\)](#) take Γ_{BIG} to be the class of game forms with perfect recall, as in [Li \(2016\)](#), and first observe that it is without loss of generality to restrict attention to Γ_1 , the class of game forms with perfect information where nature (alternatively, the mechanism’s administrator) does not randomize.

[Pycia and Troyan \(2017\)](#) also take Γ_{BIG} to be the class of game forms with perfect recall. Under preference restrictions that allow for a wide range of economic environments (such as object allocation), but rule out those with monetary transfers, they show that it is without loss of generality to restrict attention to Γ_2 , their class of *millipede* mechanisms. These mechanisms have remarkable further structure: at any information set, an agent may or may not have one *passing* action, after which he may move again; each of his other actions is a *clinching* actions, guaranteeing that the outcome belongs to a particular one of his indifference classes, after which he does not move again.

[Bade and Gonczarowski \(2017\)](#) take Γ_{BIG} to be the class of game forms for which no path from the root intersects the same information set twice; this class is larger than the class of game forms with perfect recall (because, for example, an agent may ‘forget’ the previous action he selected), but smaller than ours (because, for example, an agent may not ‘forget’ whether or not he has already selected an action). They show that it

³I am indebted to Shengwu Li for conjecturing [Theorem 5](#) to me through private correspondence.

⁴To my knowledge, [Ashlagi and Gonczarowski \(2016\)](#), [Pycia and Troyan \(2017\)](#), [Bade and Gonczarowski \(2017\)](#), and this paper were made publicly available in that order. I became aware of the other three papers after completing this paper.

is without loss of generality to restrict attention to Γ_3 , their class of *gradual revelation mechanisms*. These are our round table mechanisms with additional structure:

- no agent consecutively moves twice, or ever selects from a singleton set of actions, and
- whenever an agent’s strategy guarantees (from some history) that all possible outcomes belong to the same indifference class regardless of his preference relation, the first action of that strategy is to publicly announce his preference relation (and take no further action thereafter).

Because Γ_3 is a subset of our class of round table mechanisms, the result of [Bade and Gonczarowski \(2017\)](#) can be used to immediately strengthen our main result; we do not do so in the body of this paper in the interest of clarity about credit among recent working papers.

With respect to these other revelation principles, our contribution is not to further shrink Γ_{SMALL} , but rather to enlarge Γ_{BIG} . In particular, we drop all recall requirements and work with the largest class of game trees for which choices determine plays.

2 Base Model

2.1 Overview

A group of *agents* $N = \{1, 2, \dots, n\}$ face a set of (*public*) *outcomes* X . Each agent i has private information summarized by his *type* $\theta_i \in \Theta_i$, which determines his (complete and transitive) *preference relation* R_{θ_i} ranking outcomes. Critically, an agent’s ranking of outcomes is fixed; it would not change were he to learn anything about his peers’ types, ruling out for example auctions for mineral rights ([Milgrom and Weber, 1982](#)). A *type profile* $\theta = (\theta_i)_{i \in N}$ specifies a type for each agent, and the set of possible type profiles is given by $\Theta \equiv \times_{i \in N} \Theta_i$.

In our base model, the agents wish to condition the outcome selection on their collective information according to some *deterministic rule* $f : \Theta \rightarrow X$. Because the type profile is not common knowledge, their objective is to weakly implement f according to a particular solution concept using a game form (or mechanism).

In this section, we introduce the base model for deterministic rules and pure strategies. We begin with our class of mechanisms; describing the associated class of game trees involves a digression on order theory ([Section 2.2](#)). We then introduce the solution concept of obvious strategy-proofness for pure strategies ([Section 2.3](#)). Finally, we describe when a mechanism implements a deterministic rule, in the sense that the rule’s recommendation is plausibly reached no matter the type profile ([Section 2.4](#)). Our presentation of these notions in this section allows us to then simply describe each of our main results using an appropriate extension of the base model.

2.2 Extensive game forms

To implement the rule, we consider extensive game forms ([Osborne and Rubinstein, 1994](#); [Li, 2016](#)). Intuitively, this is an extensive-form game tree without preference information, which becomes an extensive-form game when paired with (the preference profile associated with) a type profile.

An extensive game form involves a game tree, which in its most general formulation is a collection of *histories* H that is partially ordered by *precedence* \preceq . So that our revelation principle may be compatible with as many game trees as possible, we begin by reviewing some concepts from order theory. For a fixed set H , a *partial order* \preceq is a binary relation that is reflexive, antisymmetric, and transitive, but not necessarily complete; we denote the associated strict order by \prec . For each pair $h, h' \in H$, if $h \prec h'$ and there is no h'' such that $h \prec h'' \prec h'$, then we say both that

- (i) h is an *immediate predecessor* of h' , and
- (ii) h' is an *immediate successor* of h .

A subset $H' \subseteq H$ is a *chain* if for each pair $h, h' \in H'$, either $h \preceq h'$ or $h' \preceq h$; a chain H' is *maximal* if there is no chain H'' such that $H' \subsetneq H''$; a subset $H' \subseteq H$ is *well-ordered* if for each $H'' \subseteq H'$, H'' has a minimum (in which case H' is a chain). For each pair $h, h' \in H$, the *meet of h and h'* , $h \wedge h' \in H$, is the greatest lower bound of $\{h, h'\}$, which is unique if it exists.

DEFINITION: A *meet-semilattice* is a partially ordered set (H, \preceq) such that for each pair $h, h' \in H$, there is $h \wedge h'$.

DEFINITION: A *tree* (see the [Jech, 1971](#) survey) is a partially ordered set (H, \preceq) such that for each $h \in H$, $\{h' \in H \mid h' \preceq h\}$ is well-ordered. It is *rooted* if there is $\min H$.

For our purposes, a maximal chain is a *play*, or a complete description of a sequence of choices, and a game tree is “admissible” if choices always induce a unique play. This notion is made precise in [Alós-Ferrer and Ritzberger \(2016\)](#), and we refer the reader there for the precise definition. The result that interests us here is the following:

THEOREM ([ALÓS-FERRER AND RITZBERGER, 2016](#)):⁵ A rooted “pseudotree” is “admissible” if and only if it satisfies “weak up-discreteness,” “coherence,” and “regularity.”

We use this alternative statement involving notions we find more familiar:

PROPOSITION 1: The class of meet-semilattice trees is the class of “admissible” rooted “pseudotrees.”

The proof is in [Appendix A](#). When describing these game forms, it is convenient to view the (*mechanism*) *administrator* as player 0 and write $N_0 \equiv N \cup \{0\}$ for the set of agents together with the administrator. The administrator publicly commits to how he will behave in advance, playing the same role as the ‘chance player’ or ‘nature’ in the literature on games. Following [Aumann \(1964\)](#), we include how choices can be randomized in our description; we here include only states for the administrator, and later extend the base model with different kinds of states for the agents. Formally:

DEFINITION: An *extensive game form* is given by a tuple

$$\langle H, \preceq, A, \mathcal{A}, P, (\mathcal{I}_i)_{i \in N}, g, (\Omega_h)_{h \in H_0}, (B_h)_{h \in H_0} \rangle, \text{ where}$$

⁵This is first proven in [Alós-Ferrer and Ritzberger \(2008\)](#) with a minor error, which is corrected in [Alós-Ferrer, Kern, and Ritzberger \(2011\)](#). We refer the reader to [Alós-Ferrer and Ritzberger \(2016\)](#) for this result’s most complete treatment.

1. H is the set of *histories* and \preceq is the partial order on H representing *precedence*. We require that (H, \preceq) is a meet-semilattice tree. These conditions guarantee that choices always determine a unique play, guarantee that there is a unique *initial history* h_0 which precedes all others, and allow an action to be viewed as the selection of an immediate successor. We let $\sigma(h)$ denote the immediate successors of h . A *play* is a maximal chain, which gives a complete description of a sequence of choices; we write π for a play and Π for the set of plays. A *terminal history* is a history with no successor; we write z for a terminal history and Z for the (possibly empty) set of terminal histories.
2. A is the set of *actions* and $\mathcal{A} : \cup_H \sigma(h) \rightarrow A$ is the *action function*, which at each history h associates each immediate successor $h' \in \sigma(h)$ with the action taken to reach it. We require that at any history, each available action determines a unique next history: for each $h \in H$ and each pair $h', h'' \in \sigma(h)$, $\mathcal{A}(h') \neq \mathcal{A}(h'')$. For each history h , we let $A(h) \equiv \cup_{h' \in \sigma(h)} \mathcal{A}(h')$ denote the actions available at h .
3. $P : H \setminus Z \rightarrow N_0$ is the *player function*, which associates each non-terminal history with whoever selects an action at that history—either an agent or the administrator. For each $i \in N_0$, we let $H_i \equiv \{h \in H \mid P(h) = i\}$ denote the histories that belong to i . Similarly $H_N \equiv \{h \in H \mid P(h) \in N\}$.
4. for each $i \in N$, \mathcal{I}_i is the *information partition* for i , which specifies the *information sets* partitioning H_i . We require that for each pair h, h' in the same information set I_i , the same actions are available: $A(h) = A(h')$. We write $A(I_i)$ for the actions $A(h)$ available at each history $h \in I_i$. Across all histories in a given information set I_i , i must behave the same way. Implicitly, the administrator's information set \mathcal{I}_0 consists only of singletons.
5. $g : \Pi \rightarrow X$ is the *outcome function*, which associates each play with an outcome.
6. at each $h \in H_0$, there is a set of *states* Ω_h with which the administrator can condition his action at h , and the administrator commits to $B_h : \Omega_h \rightarrow A(h)$ specifying how he will condition his act on the state's realization. In other words, the administrator commits to a behavioral strategy.

Let Γ denote the class of extensive game forms (with agents in N and outcomes in X).

2.3 Solution Concept

For the rest of this section, fix an extensive game form G . While an agent is restricted to pure strategies, he may worry that his opponents are somehow able to avoid this restriction. We therefore also introduce “choice functions” that agents are forbidden from selecting; these can be thought of as realizations of behavioral strategies that are not necessarily pure strategies.

DEFINITION: Kinds of strategies for agent i .

- a *choice function* for i is a mapping C_i which associates each history $h \in H_i$ with an available action $C(h) \in A(h)$. We use \mathcal{C}_i for the set of choice functions of i , \mathcal{C} for the set of choice function profiles, and \mathcal{C}_{-i} for the set of lists of choice functions from the agents other than i .

- a *pure strategy for i* is a choice function for i with the restriction that if h and h' share an information set, then $S(h) = S(h')$. We use \mathcal{S}_i for the set of pure strategies of i , \mathcal{S} for the set of pure strategy profiles, and \mathcal{S}_{-i} for the set of lists of pure strategies from the agents other than i .

A *solution concept* specifies a collection of strategy profiles for each game—in practice, typically the strategy profiles that are deemed ‘plausible’ in some particular sense. Here, we are only interested in games formed by pairing some extensive game form with (the preference profile associated with) a type profile: the domain is $\Gamma \times \Theta$, and each (G, θ) is associated with a collection in $2^{\mathcal{S}}$.

Our solution concept involves the repeated comparison of one’s strategy to alternatives over the course of play. To speak of strategy evaluation at a given information set, we write $\pi^G(h, C, C_0)$, or sometimes $\pi^G(h, C_i, C_{-i}, C_0)$, for the resulting play when we start from h and play proceeds according to (C, C_0) , and we write $x^G(h, C, C_0)$ for the associated outcome $g(\pi^G(h, C, C_0))$. These functions are well-defined (see [Alós-Ferrer and Ritzberger, 2016](#) and [Appendix A](#)). For each choice function $C_i \in \mathcal{C}_i$, define the *reachable histories for C_i* and the *reachable information sets for C_i* by:

$$H_i(C_i) \equiv \{h \in H_i \mid \text{there are } C_{-i}, C_0 \text{ such that } h \in \pi^G(h_\emptyset, C_i, C_{-i}, C_0)\}, \text{ and}$$

$$\mathcal{I}_i(C_i) \equiv \{I_i \in \mathcal{I}_i \mid \text{there is } h \in I_i \text{ such that } h \in H_i(C_i)\}.$$

In particular, $\mathcal{I}_i(C_i)$ is the collection of information sets where i could possibly find himself if he has been playing according to C_i , provided that any peer could realize any choice function through a behavioral strategy. We use the notation $\mathcal{I}_i(S_i)$ when the choice function is a pure strategy.

DEFINITION: Obviously dominant strategies for a fixed game (G, θ) .

- A pure strategy $S_i \in \mathcal{S}_i$ is *obviously \mathcal{S} -dominant* if for each $I_i \in \mathcal{I}_i(S_i)$, each pair $h, h' \in I_i$, each $S'_i \in \mathcal{S}_i$ such that $S_i(I_i) \neq S'_i(I_i)$,⁶ each pair $S_{-i}, S'_{-i} \in \mathcal{S}_{-i}$, and each pair $\omega_{H_0}, \omega'_{H_0} \in \Omega_{H_0}$,

$$x^G(h, S_i, S_{-i}, B_0(\omega_{H_0}))$$

$$R_i$$

$$x^G(h', S'_i, S'_{-i}, B_0(\omega'_{H_0})).$$

- A pure strategy $S_i \in \mathcal{S}_i$ is *obviously \mathcal{C} -dominant* if for each $I_i \in \mathcal{I}_i(S_i)$, each pair $h, h' \in I_i$, each $C'_i \in \mathcal{C}_i$ such that $S_i(I_i) \neq C'_i(I_i)$, each pair $C_{-i}, C'_{-i} \in \mathcal{C}_{-i}$, and each pair $C_0, C'_0 \in \mathcal{C}_0$,

$$x^G(h, S_i, C_{-i}, C_0)$$

$$R_i$$

$$x^G(h', C'_i, C'_{-i}, C'_0).$$

⁶In the original definition ([Li, 2016](#)), S_i need only be compared to S'_i if moreover S'_i can reach I_i . The two definitions are not equivalent when there is absent-mindedness, but the results in this paper hold for both definitions, essentially because any additional paths from I_i that must be considered under the given definition can never be reached and can thus be safely pruned from the game tree. We work with the given definition for its conceptual simplicity: at a given information set, the agent considers switching to any strategy that prescribes a different action.

In order for a pure strategy S_i to be obviously \mathcal{S} -dominant, it must pass a series of comparisons at each reachable information set. In particular, at each reachable information set, S_i must be compared to each pure strategy S'_i that prescribes a different action. In order for S_i to pass this comparison, each possible outcome under S_i must be at least as desirable as each possible outcome under S'_i .

The latter notion is much stronger: if a pure strategy is *obviously \mathcal{C} -dominant*, then not only is it *obviously \mathcal{S} -dominant*, but moreover, the agent entertains the possibilities that:

- he may later discover a way to coordinate across histories that share an information set,
- his opponents may later discover a way to coordinate across histories that share an information set, and
- the administrator may deviate from his announced behavioral strategy.

Under the interpretation that each player is a group with different representatives making choices at different histories (Isbell, 1957), *obvious \mathcal{C} -dominance* requires each representative to contemplate the possibility that representatives of its group, or another group, may discover a way to illegally coordinate.

Our solution concept is that each agent plays an obviously dominant strategy.

DEFINITION: Obvious strategy-proofness.

- A pure strategy profile S belongs to $\mathbf{OSP}^{\mathcal{S}}(G, \theta)$ if and only if for each agent i , S_i is obviously \mathcal{S} -dominant.
- A pure strategy profile S belongs to $\mathbf{OSP}^{\mathcal{C}}(G, \theta)$ if and only if for each agent i , S_i is obviously \mathcal{C} -dominant.

2.4 Implementation

A game form is a game with incomplete information; in particular, preference information is missing. In the context of a mechanism design model, we can say more specifically that the type profile is unknown. Following Harsanyi (1967), suppose the administrator imagines that each agent i is drawn from a population partitioned by types in Θ_i , and considers one plausible strategy for each type class. A *Harsanyi pure strategy* is a mapping $\mathbb{S}_i : \Theta_i \rightarrow \mathcal{S}_i$.

Informally, we say that a game form implements a rule if there is some Harsanyi strategy profile such that no matter how the administrator plays, for each type profile, the associated strategy profile (i) is plausible according to our solution concept, and (ii) yields the rule's recommendation. Formally:

DEFINITION: Obvious strategy-proof implementation.

- A deterministic rule f is $\mathbf{OSP}^{\mathcal{S}}$ -implementable if there are
 - (i) $G \in \Gamma$, and
 - (ii) an associated Harsanyi pure strategy profile \mathbb{S} ,
 such that for each $\theta \in \Theta$,

(i) $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^S(G, \theta)$, and

(ii) for each $\omega_{H_0} \in \Omega_{H_0}$,

$$x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) = f(\theta).$$

In this case we say G \mathbf{OSP}^S -implements f through \mathbb{S} .

- A deterministic rule f is \mathbf{OSP}^C -implementable if there are

(i) $G \in \Gamma$, and

(ii) an associated Harsanyi pure strategy profile \mathbb{S} ,

such that for each $\theta \in \Theta$,

(i) $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^C(G, \theta)$, and

(ii) for each $\omega_{H_0} \in \Omega_{H_0}$,

$$x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) = f(\theta).$$

In this case we say G \mathbf{OSP}^C -implements f through \mathbb{S} .

Of these two implementation notions, the latter is stronger.

2.5 Round table mechanisms and base revelation principle

The revelation principle for deterministic rules involves the following canonical class of mechanisms:

DEFINITION: A *round table mechanism* is a perfect-information game form such that each non-terminal history h belongs to some agent $i \in N$. Moreover, actions in $A(h)$ are subsets of Θ_i , and any two actions in $A(h)$ are pairwise disjoint. At any history of i with no strict predecessor belonging to i , the union of actions in $A(h)$ is Θ_i ; at any other history h of i , the union of actions in $A(h)$ is the intersection of the actions at all predecessors that lead to h . The class of round table mechanisms is denoted Γ^{RT} .

DEFINITION: A deterministic rule has a *round table implementation* if there is $G \in \Gamma^{RT}$ such that G \mathbf{OSP}^S -implements f through \mathbb{S} , where for each $i \in N$, each $\theta_i \in \Theta_i$, and each $h \in H_i(\mathbb{S}_i(\theta_i))$, $\mathbb{S}_i(\theta_i)(h)$ is the unique action solving

$$\theta_i \in \mathbb{S}_i(\theta_i)(h).$$

We note that \mathbf{OSP}^S -implementation with a round table mechanism implies \mathbf{OSP}^C -implementation. The logic is simple: in a round table mechanism, each choice function is a pure strategy.

Our statement of the revelation principle for deterministic rules improves on other statements by allowing all game forms for which choices determine plays and removing all recall assumptions:

PROPOSITION 2: If a deterministic rule has an \mathbf{OSP}^S -implementation, then it has a round table implementation, which is an \mathbf{OSP}^C -implementation.

The proof is in [Appendix B](#). This result is both a lemma for, and a corollary to, two of our main results: [Theorem 1](#) for implementation of stochastic rules and [Theorem 3](#) for robustness to player randomization.

3 Main Results

3.1 Revelation principle for stochastic rules

Our main result generalizes [Proposition 2](#) to the case where agents wish to condition a probabilistic distribution of outcomes on their collective information. A key step in this generalization is the elegant argument in Lemma 2 of [Pycia and Troyan \(2017\)](#) that all of the administrator's randomization can be done at the initial history; we provide this argument in our proof with some additional formality.

In order to pursue stochastic rules, we equip the set of outcomes X with a σ -algebra $\mathcal{X} \subseteq 2^X$ of measurable collections of outcomes. A *lottery* is a probability measure $p : \mathcal{X} \rightarrow [0, 1]$, and we denote the set of lotteries by \mathcal{P} . A *stochastic rule* is a function $f : \Theta \rightarrow \mathcal{P}$.

We consider only pure strategies for the implementation of stochastic rules, as issues quickly arise when the agents themselves randomize (see [Example 1](#) in [Section 3.3](#)). To extend the base model, the minimal additional structure we require is a σ -algebra \mathcal{A} on Ω_{H_0} and a probability measure $\mu : \mathcal{A} \rightarrow [0, 1]$, allowing the assignment of probabilities to measurable events upon which the administrator conditions. Let us call an extensive game form together with this additional structure (G, \mathcal{A}, μ) a *minimally-probabilistic extensive game form*, and let Γ^{MP} denote their class. We adapt our implementation notion as follows:

DEFINITION: Random obvious strategy-proof implementation.

- A stochastic rule f is *random \mathbf{OSP}^S -implementable* if there are
 - (i) $(G, \mathcal{A}, \mu) \in \Gamma^{MP}$, and
 - (ii) an associated Harsanyi pure strategy profile \mathbb{S} ,
 such that for each $\theta \in \Theta$,
 - (i) $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^S(G, \theta)$, and
 - (ii) for each $X' \in \mathcal{X}$,

$$\{\omega_{H_0} \in \Omega_{H_0} | x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) \in X'\} \in \mathcal{A},$$

$$\mu(\{\omega_{H_0} \in \Omega_{H_0} | x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) \in X'\}) = f(\theta)(X').$$

In this case we say G random \mathbf{OSP}^S -implements f through \mathbb{S} .

- A stochastic rule f is *random \mathbf{OSP}^C -implementable* if there are
 - (i) $(G, \mathcal{A}, \mu) \in \Gamma^{MP}$, and
 - (ii) an associated Harsanyi pure strategy profile \mathbb{S} ,
 such that for each $\theta \in \Theta$,
 - (i) $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^C(G, \theta)$, and
 - (ii) for each $X' \in \mathcal{X}$,

$$\{\omega_{H_0} \in \Omega_{H_0} | x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) \in X'\} \in \mathcal{A},$$

$$\mu(\{\omega_{H_0} \in \Omega_{H_0} | x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) \in X'\}) = f(\theta)(X').$$

In this case we say G random **OSP**^C-implements f through \mathbb{S} .

We note that it is sometimes possible to represent the probabilities specified by μ as calculations involving independent probabilities across different histories in H_0 , for example when each of the administrator's state spaces is the unit interval together with the Lebesgue measurable sets and the Lebesgue measure (Aumann, 1964). Our approach does not require any such restrictions, and even allows the administrator to correlate his choices across histories. Finally, we remark that as long as values of μ can be assigned to the events specified above, a suitable \mathcal{A} can be recovered.⁷

We emphasize that this notion of random implementation does not require lotteries to be assigned to terminal nodes. Given such a game form, one could replace each terminal node with a move by the administrator where he randomly selects an outcome according to the associated lottery. One could also consider game forms where the administrator randomizes at various histories scattered throughout the game tree. In each of these cases, it is possible for a stochastic rule to be implemented.

In our canonical class of mechanisms, the administrator moves once and only once, beginning play by selecting a round table mechanism for the agents:

DEFINITION: A *randomized round table mechanism* is a perfect-information minimally-probabilistic extensive game form such that the initial history belongs to the administrator 0, and every other non-terminal history h belongs to some agent $i \in N$. Moreover, at each history h of player i , actions in $A(h)$ are subsets of Θ_i , and any two actions in $A(h)$ are pairwise disjoint. At any history of i with no strict predecessor belonging to i , the union of actions in $A(h)$ is Θ_i ; at any other history h of i , the union of actions in $A(h)$ is the intersection of the actions at all predecessors that lead to h . The class of randomized round table mechanisms is denoted Γ^{RRT} .

DEFINITION: A stochastic rule has a *randomized round table implementation* if there is $(G, \mathcal{A}, \mu) \in \Gamma^{RRT}$ such that (G, \mathcal{A}, μ) **OSP**^S-implements f through \mathbb{S} , where for each $i \in N$, each $\theta_i \in \Theta_i$, and each $h \in H_i(\mathbb{S}_i(\theta_i))$, $\mathbb{S}_i(\theta_i)(h)$ is the unique action solving

$$\theta_i \in \mathbb{S}_i(\theta_i)(h).$$

We note that random **OSP**^S-implementation with a randomized round table mechanism implies random **OSP**^C-implementation: each choice function is a pure strategy, and the administrator is done playing at each player history. Our main result is that it is safe to restrict attention to this class:

THEOREM 1: If a stochastic rule has a random **OSP**^S-implementation, then it has a randomized round table implementation, which is a random **OSP**^C-implementation.

⁷To see this, for each $\theta \in \Theta$, consider the mapping $x_\theta^G : \Omega_{H_0} \rightarrow X$ that takes each state profile ω_{H_0} to the outcome $x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0}))$. Because \mathcal{X} is a σ -algebra, it is well known that the inverse images of x_θ^G form a σ -algebra \mathcal{A}_θ . Moreover, it is well-known that an arbitrary intersection of σ -algebras is itself a σ -algebra; thus \mathcal{A} may be taken to be $\cap \mathcal{A}_\theta$. It is straightforward to complete the measure μ from the given probabilities.

The proof is in [Appendix C](#). In the next section, we establish that this generalization is necessary, in the sense that it is *not* safe to restrict attention to the canonical class for deterministic rules.

3.2 Insufficiency of round table mechanisms for stochastic rules

In this section, we demonstrate by means of a simple example that [Proposition 2](#) should not be used when investigating the implementation of stochastic rules.

There is a set of agents $N = \{1, 2, 3\}$ with generic members i, j, k and a set of objects $\{A, B, C\}$ with generic members a, b, c . An *allocation* is a bijection between the agents and objects, associating each agent with his own object, and an outcome in X is a lottery over allocations. For each $i \in N$, Θ_i is the set of von Neumann-Morgenstern utility representations of preferences over lotteries of objects; as there is no indifference, we normalize so that for each θ_i , there is some m satisfying $9 > m > 0$ such that the range of θ_i is $\{0, m, 9\}$.

DEFINITION: For each strict order of agents i then j then k , the rule $f^{i,j,k}$ gives i his favorite object, j his favorite of the other two objects, and k the remaining object. The *random serial dictatorship* f is given by the uniform randomization over these rules:

$$f \equiv \frac{1}{6} \sum f^{i,j,k}.$$

THEOREM 2: Random serial dictatorship has no round table implementation.

The proof is in [Appendix D](#). When paired with the observation by [Pycia and Troyan \(2017\)](#) that random serial dictatorship does have a randomized round table implementation, [Theorem 2](#) establishes that [Theorem 1](#) is a necessary generalization of [Proposition 2](#) for the implementation of stochastic rules.

3.3 Robustness to player randomization

Does our revelation principle hold if, as assumed throughout most of game theory, the players might randomize? The answer depends on how this question is made precise, which involves a series of subtle modeling choices. We proceed briefly through the same steps from [Section 3](#), occasionally pausing to highlight a choice of conceptual interest.

First of all, how should we model player randomization? Though for game forms with perfect recall, behavioral strategies are equivalent to mixed strategies ([Kuhn, 1953](#); [Aumann, 1964](#)), this is not true for all game forms we consider; we therefore consider both. For our result, we are able to accommodate unusually general notions of randomization that do not require probabilities:

- For mixed strategies, let us say a *mixed extensive game form* $(G, (\Omega_i)_{i \in N})$ consists of (i) an extensive game form G , and (ii) for each $i \in N$, a set of *states* Ω_i with which the player can condition his pure strategy in a mixed strategy. A *mixed strategy for i* is a mapping $M_i : \Omega_i \rightarrow \mathcal{S}_i$ specifying how i conditions his selection of pure strategy on the realization of a state at the start of the game. We use \mathcal{M}_i , \mathcal{M} , and \mathcal{M}_{-i} to denote, respectively, the set of mixed strategies of i , the set of mixed strategy profiles, and the set of lists of mixed strategies from agents other than i .

- For behavioral strategies, let us say a *behavioral extensive game form* $(G, (\Omega_h)_{h \in H_N})$ consists of (i) an extensive game form G , and (ii) at each $h \in H_N$, a set of *states* Ω_h with which the player at h can condition his action at h in a behavioral strategy. We require that for each pair h, h' in the same information set, $\Omega_h = \Omega_{h'}$; each is a copy of the other, though their realizations may be different. Thus at a given information set, the agent cannot infer anything about his history simply by looking at his state space. A *behavioral strategy for i* is a list of mappings $B = (B_h)_{h \in H_i}$, where $B_h : \Omega_h \rightarrow A(h)$ specifies how i conditions his selection of an action at h on the realization of a state at h , with the restriction that if h and h' share an information set, then $B_h = B_{h'}$. We use \mathcal{B}_i , \mathcal{B} , and \mathcal{B}_{-i} to denote, respectively, the set of mixed strategies of i , the set of mixed strategy profiles, and the set of lists of mixed strategies from agents other than i .

While we do not require randomization to involve probabilities, the addition of further familiar structure would of course not impact our results.⁸

In order to specify where an agent assesses a given strategy, we choose notions of reachable information sets $\mathcal{I}_i(M_i)$ and $\mathcal{I}_i(B_i)$ that are analogous to our original definition, simply modified to allow any realization of agent i 's states. The more interesting question is *when* an agent assesses his strategy at a given information set: before or after states are realized?

There is a subtle conceptual problem with taking an *ex-ante* approach, which is best illustrated with behavioral strategies. At a given information set, suppose B_i and B'_i prescribe the same range of outcomes. Before the state is realized, there is a straightforward process for determining whether or not the two are equivalent: for each state, check whether or not B_i and B'_i prescribe the same action; if they agree at every state, they are equivalent; otherwise, they are distinct. The problem is that this process involves state-by-state *contingent reasoning*, and one of the primary motivations for obvious strategy-proofness is the observation that subjects may have difficulty with such reasoning (Li, 2016). An agent without this ability who is tasked with making comparisons *before* the state is realized would therefore either have to (i) compare a strategy to itself or (ii) never compare a strategy to another that prescribes the same range of actions. We believe that neither of these alternatives is in the spirit of the original notion, and therefore use the following *ex-post* notions:

- For *obvious \mathcal{M} -dominance*, we assume that before the game begins, the state ω_i is revealed to the agent, and this information is recorded and available thereafter. At any given information set, then, each mixed strategy is already reduced to a pure strategy.
- For *obvious \mathcal{B} -dominance*, we do not assume that the agent has a record of every state that is revealed to him over the course of the game; indeed, without perfect recall, the agent may not even have a record of every action that he has taken. At

⁸For example, we could require for each state space: (i) a collection of events whose relative likelihood can be compared; (ii) that this collection is a σ -algebra; (iii) that each agent's likelihood comparisons have representation by a probability measure; and (iv) that these probability measures all correspond to some 'objective' measure. To guarantee that conditional probabilities can be defined and used to calculate lotteries of outcomes, we could impose that these probability spaces are *standard* (Mackey, 1957), provide σ -algebras for acts and outcomes, and impose measurability restrictions on mappings (Halmos, 1950; Aumann, 1964). To rule out strategy correlation (Aumann, 1987), we could impose that the product state spaces are associated with their product σ -algebras and product measures.

a given information set, however, the current state ω_x is revealed, and the agent takes this information into account when making comparisons—though he does not know which of his information set’s histories he is playing from, and therefore does not know which of these histories ω_x belongs to.

For brevity, we omit the formal definitions of *obvious \mathcal{M} -dominance* and *obvious \mathcal{B} -dominance*; these notions should be clear given this discussion and the analogous definition for pure strategies. The solution concepts $\mathbf{OSP}^{\mathcal{M}}(G, \theta)$ and $\mathbf{OSP}^{\mathcal{B}}(G, \theta)$, for mixed and behavioral strategies, respectively, are also analogous to the definition for pure strategies.

We use $\mathbb{M}_i : \theta_i \rightarrow \mathcal{M}_i$ for a *Harsanyi mixed strategy* and $\mathbb{B}_i : \theta_i \rightarrow \mathcal{B}_i$ for a *Harsanyi behavioral strategy*. Informally, we say that a game form implements a rule if there is some Harsanyi strategy profile such that, for each type profile, the associated strategy profile (i) is plausible according to our solution concept, and (ii) yields the rule’s recommendation. There is one last subtle point that still must be addressed, however. Suppose the mechanism administrator wishes to implement a deterministic rule, yet he and the agents may randomize. When should we say he has succeeded? There are two natural definitions:

1. for each $\theta \in \Theta$, the outcome is $f(\theta)$ for each joint realization of states, or
2. for each $\theta \in \Theta$, the outcome is $f(\theta)$ with probability 1.

Under the second notion, we are unable pursue our approach for our revelation principle, as we cannot replace any randomized strategy with any of its realizations:

EXAMPLE 1: There is one agent with one possible type, and there are two outcomes, x and x' . The agent is indifferent between these outcomes, and the rule specifies that x should be selected. In the extensive game form G , the agent first selects a number from the unit interval, then does so once more, at which point the game ends. Moreover, the agent is absent-minded; in particular, all histories share an information set. The outcome function assigns x' only when the same number is selected twice, and assigns x otherwise. The behavioral strategy B selects a number from the unit interval using the uniform distribution at the sole information set. Though (G, B) yields outcome x with probability 1, for each pure strategy S and each mixed strategy M , (G, S) and (G, M) both yield outcome x' with probability 1.

We therefore use the first notion, which is stronger and allows us to use a less-structured model. Thus our implementation notion is not that the rule’s recommendation is selected almost surely, but rather that it is selected surely—for every realization of states. For brevity, we omit the formal definitions for a rule to be $\mathbf{OSP}^{\mathcal{M}}$ -implementable and $\mathbf{OSP}^{\mathcal{B}}$ -implementable; these notions should be clear given this discussion and the analogous definition for pure strategies.

Altogether, we find that for deterministic rules, our revelation principle is robust to player randomization:

THEOREM 3: If a deterministic rule f is either

- (i) $\mathbf{OSP}^{\mathcal{S}}$ -implementable,
- (ii) $\mathbf{OSP}^{\mathcal{M}}$ -implementable, or
- (iii) $\mathbf{OSP}^{\mathcal{B}}$ -implementable,

then it has a round table implementation, which is an OSP^C -implementation.

The proof is in [Appendix E](#). On the one hand, we find that any deterministic rule that can be implemented by taking advantage of agent randomization can also be implemented with each agent playing a pure strategy, which is obviously dominant to him even if he worries that the other agents might randomize. On the other hand, it is unclear that this is true for stochastic rules. Altogether, then, while it is possible to reassure agents that they need not randomize or worry about other agents randomizing, it is difficult to rely on the agents to provide the desired randomization themselves; this should be provided by the administrator.

3.4 Justification for ordinal mechanisms

In this section, we state a result inspired by a recent paper ([Carroll, 2017](#)), who seeks a justification for stochastic rules that elicit only ordinal information from the agents about their preferences over outcomes.

Suppose that we have a finite set of outcomes, and each agent's type space consists of von Neumann-Morgenstern utility functions representing preferences over lotteries. Moreover, each agent's type is compatible with a strict ranking over outcomes. Let us say that a stochastic rule is *ordinal* if for each $i \in N$, each $\theta_{-i} \in \Theta_{-i}$, and each pair $\theta_i, \theta'_i \in \Theta_i$ compatible with the same ranking, $f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i})$. In other words, a stochastic rule is ordinal if it only depends on the agents' rankings of deterministic outcomes. Finally, let a *randomized ordinal round table implementation* be a randomized round table implementation when each agent's type space is replaced by the associated space of rankings.

THEOREM 4: Suppose there is a finite set of outcomes and each agent has a von Neumann-Morgenstern type compatible with a strict ranking over outcomes. If a stochastic rule f has a random OSP^S -implementation, then f is ordinal and has a randomized ordinal round table implementation.

The proof is in [Appendix F](#). In this sense, obvious strategy-proofness provides a foundation for ordinal mechanisms. This observation was inspired by [Carroll \(2017\)](#), who provides a foundation involving the *interdependence* of the agents' preferences over lotteries. The alternative foundation offered here does not rely on interdependence as it involves a stronger solution concept.

3.5 Round table mechanisms and ordinary strategy-proofness

It turns out to be fruitful to consider standard *strategy-proofness* in the context of round table mechanisms, a point I am grateful to Shengwu Li for suggesting to me. For brevity, we gather the necessary definitions, which were introduced in sequence for obvious strategy-proofness, here:

DEFINITION: Strategy-proofness.

- A pure strategy $S_i \in \mathcal{S}_i$ is *dominant* if for each $S'_i \in \mathcal{S}_i$, each $S_{-i} \in \mathcal{S}_{-i}$, and each $\omega_{H_0} \in \Omega_{H_0}$,

$$\begin{aligned}
& x^G(h_\emptyset, S_i, S_{-i}, B_0(\omega_{H_0})) \\
& \quad R_i \\
& x^G(h_\emptyset, S'_i, S_{-i}, B_0(\omega_{H_0})).
\end{aligned}$$

- A pure strategy profile S belongs to $\mathbf{SP}(G, \theta)$ if and only if for each agent i , S_i is dominant.
- A deterministic rule f is \mathbf{SP} -implementable if there are

(i) $G \in \Gamma$, and

(ii) an associated Harsanyi pure strategy profile \mathbb{S} ,

such that for each $\theta \in \Theta$,

(i) $(\mathbb{S}_i(\theta_i)) \in \mathbf{SP}(G, \theta)$, and

(ii) for each $\omega_{H_0} \in \Omega_{H_0}$,

$$x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) = f(\theta).$$

In this case we say G \mathbf{SP} -implements f through \mathbb{S} .

Which mechanisms \mathbf{SP} -implement a given rule? It could be that there are none, but if there are any, then by the standard revelation principle (Gibbard, 1973; Myerson, 1981), at least one of them is a direct mechanism. Beyond this, very little is known in general, but for our purposes, all that matters is whether or not one of them is a round table mechanism:

THEOREM 5: A deterministic rule f has an \mathbf{OSP}^c -implementation if and only if it has a \mathbf{SP} -implementation through a round table mechanism.

See Figure 1; the proof is in Appendix G. Thus the analysis of whether a deterministic rule has an obviously strategy-proof implementation may be safely confined to a simple class of mechanisms and a familiar solution concept.

Appendix A

In this appendix, we prove Proposition 1. We begin by summarizing some known results about game trees. As we do not work directly with conditions in marked by quotation marks, and the reader is referred to the cited sources for their definitions. By contrast, we do work with the italicized conditions, which are defined after our summary.

We wish for our extensive game forms to be compatible with a large class of game trees, subject to the “admissibility” requirement that any choices determine a unique play. Fortunately, Alós-Ferrer and Ritzberger (2008) (henceforth, AFR) have provided (independent) necessary and sufficient conditions:

THEOREM (AFR, THEOREM 6):⁹ A rooted pseudotree is “admissible” if and only if it satisfies “weak up-discreteness,” “coherence,” and *regularity*.

⁹A correction to the definition of “admissible” is provided in Alós-Ferrer, Kern, and Ritzberger (2011), and the most complete treatment of this result is given in Alós-Ferrer and Ritzberger (2016).




Figure 1: Theorem 5. For a given deterministic rule f , let Γ_f be the class of game forms that **SP**-implement f . In general, Γ_f may be empty (not pictured). Otherwise, by the classic revelation principle, Γ_f intersects the class of direct mechanisms. If Γ_f does not intersect the class of round table mechanisms (left), then f has no **OSP**-implementation. On the other hand, if Γ_f does intersect the class of round table mechanisms (right), then those mechanisms in the intersection are **OSP**-implementations.

AFR work with set-based rooted pseudotrees, where each history is the collection of plays that include it and precedence is set inclusion; this is without loss of generality in a precise sense (Alós-Ferrer and Ritzberger, 2005b). Because “weak up-discreteness” and “coherence” are together equivalent to *up-discreteness* (AFR, Corollary 3), “admissibility” is equivalent to *up-discreteness* and *regularity* (AFR, Corollary 5). Moreover, *up-discreteness* is equivalent to being a tree (Lemma AFR 1).

Furthermore, “admissibility” trivially implies “coherence” and *regularity*, which imply “selectiveness” (AFR, Proposition 6b).¹⁰ “Selectiveness” implies “well-meeting” (AFR, Footnote 27),¹¹ which is equivalent to being a meet-semilattice (Alós-Ferrer and Ritzberger, 2005a). Additionally, “selectiveness” implies *regularity* (FR, Proposition 6a).

These existing results together imply a *regular* rooted tree is a meet-semilattice (Lemma AFR 2). To assure the reader that this is indeed true in general for order-theoretic pseudotrees, we provide a direct proof. We then make two novel observations: a meet-semilattice pseudotree is *regular* (Lemma 1.1), and a meet-semilattice tree is rooted (Lemma 1.2). It is immediate from the lemmas that “admissibility” is equivalent to being a (necessarily rooted) meet-semilattice tree.

Throughout, we assume the Axiom of Choice, which is equivalent to the Hausdorff Maximality Principle, which is equivalent to Zorn’s Lemma. The Axiom of Choice guarantees that strategies exist.

For each $h \in H$, let $\uparrow h \equiv \{h' \in H \mid h' \preceq h\}$ and $\uparrow\uparrow h \equiv (\uparrow h \setminus \{h\})$ denote the weak and strict predecessors of h , respectively. We restate the definitions from the text using this notation:

DEFINITION: A *pseudotree* (Koppelberg and Monk, 1991) is a partially ordered set (H, \preceq) such that for each $h \in H$, $\uparrow h$ is a chain. It is *rooted* if there is $\min H$.

¹⁰In fact, they show that *regularity* and a condition weaker than “coherence” imply “selectiveness”.

¹¹In this article, earlier histories are *smaller*, while in AFR, earlier histories are *greater*; thus we use “well-meeting” in place of their “well-joining.”

DEFINITION: A *tree* (see the [Jech, 1971](#) survey) is a partially ordered set (H, \lesssim) such that for each $h \in H$, $\uparrow h$ is well-ordered.

DEFINITION: A pseudotree (H, \lesssim) is *up-discrete* if each nonempty chain has a minimum.

The following observation appears in [Alós-Ferrer and Ritzberger \(2008\)](#); we provide the simple proof for completeness:

LEMMA AFR 1: A pseudotree is a tree if and only if it is *up-discrete*.

PROOF: We prove the two implications in sequence.

[\Rightarrow] Let H' be a nonempty chain. Then there is $h \in H'$. Since (H, \lesssim) is tree, there is $h_m \equiv \uparrow h$. Let $h' \in H'$. Since H' is a chain, either $h' \lesssim h$ or $h \prec h'$. In the former case, $h' \in \uparrow h$, so $h_m \lesssim h'$ by construction; in the latter case, $h_m \lesssim h \prec h'$. Since $h' \in H'$ was arbitrary, thus $h_m = \min H'$.

[\Leftarrow] Let $h \in H$ and let $H' \subseteq \uparrow h$, $H' \neq \emptyset$. Since (H, \lesssim) is a pseudotree, thus $\uparrow h$ is a chain, so H' is a nonempty chain. By *up-discreteness*, H' has a minimum. Since H' was arbitrary, thus $\uparrow h$ is well-ordered. ■

Before proceeding, we introduce some more notation. For each $H^* \subseteq H$, let $\mathcal{U}(H^*)$ denote the set of upper bounds of H^* . Recall that $\sup H^* = \min \mathcal{U}(H^*)$; since \lesssim is a partial order, if $\sup H^*$ exists then it is unique. Similarly, let $\mathcal{L}(H^*)$ denote the set of lower bounds of H^* , and recall that $\inf H^* = \max \mathcal{L}(H^*)$, which is unique if it exists. Finally, we write $h \wedge h'$ for the *meet of h and h'* , $\inf\{h, h'\}$, which is unique if it exists.

DEFINITION: A *meet-semilattice* is a partially ordered set (H, \lesssim) such that for each pair $h, h' \in H$, there is $h \wedge h'$.

DEFINITION: A pseudotree (H, \lesssim) is *regular* if for each $h \in H$ such that $(\uparrow\uparrow h) \neq \emptyset$, there is $\sup(\uparrow\uparrow h)$.

We now prove the lemmas in sequence.

LEMMA AFR 2: If a rooted tree is *regular*, then it is a meet-semilattice.

PROOF: Let $h, h' \in H$. If h and h' can be compared, then trivially there is $h \wedge h' \in \{h, h'\}$, so assume h and h' cannot be compared.

Since (H, \lesssim) is a tree, both $\uparrow h$ and $\uparrow h'$ are well-ordered. Since $h \in (\uparrow h \setminus \uparrow h')$ and $h' \in (\uparrow h' \setminus \uparrow h)$, thus there are $h_m \equiv \min(\uparrow h \setminus \uparrow h')$ and $h'_m \equiv \min(\uparrow h' \setminus \uparrow h)$. Notice that h_m and h'_m cannot be compared: $h_m \lesssim h'_m$ implies $h_m \in \uparrow h'$, while $h'_m \lesssim h_m$ implies $h'_m \in \uparrow h$.

Define $H^* \equiv (\uparrow\uparrow h_m)$. Since $h_m \not\lesssim h'_m$, thus h_m is not the root, so H^* is not empty as it includes the root. We claim $H^* = \uparrow h \cap \uparrow h'$.

To see $H^* \subseteq \uparrow h \cap \uparrow h'$, let $h^* \in H^*$. Then $h^* \prec h_m \lesssim h$, so $h^* \in \uparrow h$. Since $h^* \prec h_m = \min(\uparrow h \setminus \uparrow h')$, thus $h^* \in \uparrow h \cap \uparrow h'$.

Conversely, let $h^* \in \uparrow h \cap \uparrow h'$. Then $h^* \notin \{h_m, h'_m\}$. Since $\uparrow h$ and $\uparrow h'$ are chains, thus h^* can be compared to h_m and h^* can be compared to h'_m . But we cannot have $h_m \prec h^*$, else h_m and h'_m can be compared: if $h^* \prec h'_m$, by transitivity; if $h'_m \prec h^*$, because $\uparrow h^*$ is a chain as (H, \lesssim) is a tree. Thus $h^* \prec h_m$, so $h^* \in H^*$.

Altogether, $H^* = \uparrow h \cap \uparrow h'$, as desired. We next claim H^* has a maximum. Indeed, assume by way of contradiction that H^* has no maximum. Since $H^* = (\uparrow\uparrow h_m)$ is nonempty, by *regularity* it has a supremum, so $\mathcal{U}(H^*)$ has a minimum. Since $h_m \in \mathcal{U}(H^*)$, and since $h^* \prec h_m$ implies $h^* \in H^*$ which implies $h^* \notin \mathcal{U}(H^*)$, thus $h_m = \min \mathcal{U}(H^*)$. But $H^* \subseteq \uparrow h'$, so $h' \in \mathcal{U}(H^*)$ and $h_m \not\lesssim h'$, contradicting that $h_m = \min \mathcal{U}(H^*)$. Thus H^* has maximum h_M^* , as desired.

Finally, we claim $h_M^* = h \wedge h'$. Indeed, since $h_M^* = \max \uparrow h \cap \uparrow h'$, thus $h_M^* \in \mathcal{L}(\{h, h'\})$. Assume, by way of contradiction, there is $h^* \in \mathcal{L}(\{h, h'\})$ such that $h_M^* \prec h^*$. Since $\uparrow h$ is a chain, thus h^* and h_m can be compared. Moreover, $h_M^* = \max \uparrow\uparrow h_m$ and $h_M^* \prec h^*$, so $h_m \lesssim h^*$. But then $h_m \lesssim h^* \lesssim h'$, so $h_m \in \uparrow h'$, contradicting $h_m = \min \uparrow h \setminus \uparrow h'$. Thus $h_M^* = \max \mathcal{L}(\{h, h'\}) = h \wedge h'$, as desired. ■

LEMMA 1.1: If a pseudotree is a meet-semilattice, then it is *regular*.

PROOF: Assume, by way of contradiction, (H, \lesssim) is not *regular*. Then there is $h \in H$ such that $\uparrow\uparrow h$ is nonempty with no supremum.

Assume, by way of contradiction, $\uparrow\uparrow h \cap \mathcal{U}(\uparrow\uparrow h) \neq \emptyset$. Since (H, \lesssim) is a pseudotree, thus $\uparrow\uparrow h$ is chain. Then for each pair $h', h'' \in \uparrow\uparrow h \cap \mathcal{U}(\uparrow\uparrow h)$, since $h' \not\lesssim h''$ and $h'' \not\lesssim h'$, thus $h' \sim h''$, so $h' = h''$. But then $\uparrow\uparrow h \cap \mathcal{U}(\uparrow\uparrow h)$ is a singleton, so there is $\min \uparrow\uparrow h \cap \mathcal{U}(\uparrow\uparrow h) = \max \uparrow\uparrow h = \sup \uparrow\uparrow h$, contradicting that there is no $\sup \uparrow\uparrow h$. Thus $\uparrow\uparrow h \cap \mathcal{U}(\uparrow\uparrow h) = \emptyset$, as desired.

By construction, $h \in \mathcal{U}(\uparrow\uparrow h)$. Moreover, for each $h' \in H$, $h' \prec h$ implies $h' \in \uparrow\uparrow h$, which implies $h' \notin \mathcal{U}(\uparrow\uparrow h)$. Since $h \neq \min \mathcal{U}(\uparrow\uparrow h)$, and since $h' \in \mathcal{U}(\uparrow\uparrow h)$ implies $h' \not\prec h$, thus there is $h' \in \mathcal{U}(\uparrow\uparrow h)$ such that h and h' cannot be compared.

Since (H, \lesssim) is a meet-semilattice, there is $h^* \equiv h \wedge h'$. Since h and h' cannot be compared, thus $h^* \notin \{h, h'\}$, so $h^* \in \uparrow\uparrow h$. Since $\uparrow\uparrow h$ is a chain with no maximum, there is $h^{**} \in \uparrow\uparrow h$ such that $h^* \prec h^{**}$. But then since $h, h' \in \mathcal{U}(\uparrow\uparrow h)$, thus $h^{**} \in \mathcal{L}(\{h, h'\})$ and $h^* \prec h^{**}$, contradicting that $h^* = h \wedge h'$. Thus (H, \lesssim) is regular, as desired. ■

LEMMA 1.2: Every meet-semilattice tree is rooted.

PROOF: Let (H, \lesssim) satisfy the hypotheses. By Lemma AFR 1, (H, \lesssim) is *up-discrete*, so each nonempty chain has a lower bound. Thus by Zorn's Lemma, H has minimal element h_\emptyset .

Assume, by way of contradiction, H has two distinct minimal elements h_\emptyset and h'_\emptyset . By *meet-closure*, there is $h^* \equiv h_\emptyset \wedge h'_\emptyset$. Since h_\emptyset and h'_\emptyset cannot be compared, thus $h^* \notin \{h_\emptyset, h'_\emptyset\}$. But then $h^* \prec h_\emptyset$, contradicting the minimality of h_\emptyset . Thus $h_\emptyset = \min H$. ■

Altogether, we have established:

PROPOSITION 1: The class of meet-semilattice trees is the class of “admissible” rooted pseudotrees.

Appendix B

In this appendix, we prove [Proposition 2](#). The overall approach is to take an arbitrary deterministic rule f and assume that some game form G $\mathbf{OSP}^{\mathcal{S}}$ -implements f through some Harsanyi pure profile \mathbb{S} , then successively transform the pair (G, \mathbb{S}) until we have a round table implementation of f .

We first remove all of the administrator's moves; this involves 'pruning' off irrelevant branches as in [Li \(2016\)](#), as well as 'squishing' nodes that have a single immediate successor. Formally:

DEFINITION: Tree operations. Let (H, \preceq) be a meet-semilattice tree, let $H' \subseteq H$ be nonempty, and let \preceq' be the restriction of \preceq to H' . Then

- (H', \preceq') is a *pruning* of (H, \preceq) if there is $E \subseteq H$ such that

$$H' = H \setminus \{h \in H \mid \text{there is } e \in E \text{ such that } e \preceq h\}, \text{ and}$$

- (H', \preceq') is a *squishing* of (H, \preceq) if there is $H^* \subseteq H$ such that
 - for each $h^* \in H^*$, h^* has one immediate successor in H , and
 - $H' = H \setminus H^*$.

Because we use a particularly large game trees, one might worry that these operations might result in an illegal partially ordered set. Fortunately, however, this is not the case.

First, we prove that pruning a meet-semilattice tree produces a meet-semilattice tree:

LEMMA 2.1: If (H, \preceq) is a meet-semilattice tree and (H', \preceq') is a pruning of (H, \preceq) , then (H', \preceq') is a meet-semilattice tree.

PROOF: Let (H, \preceq) and (H', \preceq') be as in the hypothesis. Then there is $E \subseteq H$ such that $H' = H \setminus \{h \in H \mid \text{there is } e \in E \text{ such that } e \preceq h\}$.

We first claim that for each $h' \in H'$, $\{h \in H' \mid h \preceq h'\} = \{h \in H \mid h \preceq h'\}$. Indeed, let $h' \in H'$. By construction, $\{h \in H' \mid h \preceq h'\} \subseteq \{h \in H \mid h \preceq h'\}$. Conversely, let $h \in H$ such that $h \preceq h'$, and assume, by way of contradiction, $h \notin H'$. Then there is $e \in E$ such that $e \preceq h$. But then by transitivity, $e \preceq h'$, so $h' \notin H'$, contradicting $h' \in H'$. Altogether, $\{h \in H' \mid h \preceq h'\} = \{h \in H \mid h \preceq h'\}$, as desired.

To see (H', \preceq') is a tree, let $h' \in H'$. By the claim, h' has the same predecessors in H and H' . Moreover, this set is well-ordered since (H, \preceq) is a tree. Since $h' \in H'$ was arbitrary, we are done.

To see (H', \preceq') is a meet-semilattice, let $h, h' \in H'$. Since h has the same predecessors in H and H' , and since h' does as well, thus h and h' have the same lower bounds in H and H' . Moreover, this set has a maximum since (H, \preceq) is a meet-semilattice. Since $h, h' \in H'$ were arbitrary, we are done. ■

Squishing a meet-semilattice tree produces a meet-semilattice tree, as well:

LEMMA 2.2: If (H, \preceq) is a meet-semilattice tree and (H', \preceq') is a squishing of (H, \preceq) , then (H', \preceq') is a meet-semilattice tree.

PROOF: Let (H, \lesssim) and (H', \lesssim') be as in the hypothesis. Then there is $H^* \subseteq H$ such that for each $h^* \in H^*$, h^* has one immediate successor in H , and (ii) $H' = H \setminus H^*$.

To see (H', \lesssim') is a tree, let $h' \in H'$ and let $H^* \subseteq \{h \in H' \mid h \lesssim h'\}$. Then $H^* \subseteq \{h \in H \mid h \lesssim h'\}$. Since (H, \lesssim) is a tree, $\{h \in H \mid h \lesssim h'\}$ is well-ordered, so H^* has a minimum. Since $H^* \subseteq \{h \in H' \mid h \lesssim h'\}$ was arbitrary, thus $\{h \in H' \mid h \lesssim h'\}$ is well-ordered. Since $h' \in H'$ was arbitrary, we are done.

Next, we claim that for each pair $h, h^* \in H$, $h^* \prec h$ implies h^* has an immediate successor $h^{**} \in H$ such that $h^{**} \lesssim h$. Indeed, let $h, h^* \in H$ with $h^* \prec h$. Since (H, \lesssim) is a tree and $\{h' \in H \mid h^* \prec h' \lesssim h\}$ is a nonempty subset of $\{h' \in H \mid h' \lesssim h\}$, thus $\{h' \in H \mid h^* \prec h' \lesssim h\}$ has minimum h^{**} . There cannot be h^{***} such that $h^* \prec h^{***} \prec h^{**}$, else we would have $h^{***} \in \{h' \in H \mid h^* \prec h' \lesssim h\}$, contradicting the minimality of h^{**} . Thus $h^{**} \in H$ is an immediate successor of $h^* \in H$ such that $h^{**} \lesssim h$, as desired.

Assume, by way of contradiction, (H', \lesssim') is not a meet-semilattice. Then there are $h, h' \in H'$ with no meet in H' . Since (H, \lesssim) is a meet-semilattice, h and h' have meet h^* in H . If h^* were in H' , then it would be the meet of h and h' in H' , which does not exist; thus $h^* \in H \setminus H'$. In particular, $h^* \notin \{h, h'\}$, so $h^* \prec h$ and $h^* \prec h'$. By the claim, h^* has immediate successors h^{**} and h^{***} such that $h^{**} \lesssim h$ and $h^{***} \lesssim h'$. But $h^* \in H \setminus H'$, so h^* has a unique immediate successor, so $h^{**} = h^{***}$, so h^{**} is a lower bound of h and h' that is greater than h^* , contradicting that h^* is the meet of h and h' . ■

At this point, we proceed to take pairs (G, \mathbb{S}) and transform them into pairs $(\mathcal{T}(G), \mathcal{T}(\mathbb{S}))$ according to a series of transformations \mathcal{T} while preserving implementation. As promised, the first transformation removes the administrator from the game form:

DEFINITION: Transformation $\mathcal{T}_{\omega_{H_0}}$.

- For each $\omega_{H_0} \in \Omega_{H_0}$, define $\mathcal{T}_{\omega_{H_0}}(G)$ by $G' = \langle H', \lesssim', A', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g' \rangle$:
 - (i) $C_0 \equiv B_0(\omega_{H_0})$,
 - (ii) $E \equiv \cup_{h \in H_0} \{h' \in \sigma(h) \mid \mathcal{A}(h') \neq C_0(h)\}$
 - (iii) $H^* \equiv H \setminus \{h \in H \mid \text{there is } e \in E \text{ such that } e \lesssim h\}$,
 - (iv) $H' \equiv H^* \setminus H_0$, and
 - (v) everything else is the associated restriction.
- Given \mathbb{S} for G , let $\mathcal{T}_{\omega_{H_0}}(\mathbb{S})$ be the associated restriction for $\mathcal{T}_{\omega_{H_0}}(G)$.

LEMMA 2.3: If $G \in \Gamma$ \mathbf{OSP}^S -implements f through \mathbb{S} , then for each $\omega_{H_0} \in \Omega_{H_0}$, $\mathcal{T}_{\omega_{H_0}}(G)$ \mathbf{OSP}^S -implements f through $\mathcal{T}_{\omega_{H_0}}(\mathbb{S})$.

We omit the proof as the notation obscures the logic. For a fixed realization of the administrator's states, the transformation first removes all successors of any node that the administrator passes over, then squishes the administrator's nodes. This is a composition of squishing with pruning, so the resulting game tree is admissible. In the transformed game form, whenever an agent i evaluates a given strategy S_i , there are fewer possibilities; thus the worst-case scenario for S_i can only improve, while the best-case scenario across alternatives can only worsen. Because this is true across agents and type profiles, \mathbf{OSP}^S -implementation is preserved.

With the administrator removed, we next remove absent-mindedness. For a given history h , let $\mathcal{E}(h)$ denote the earliest predecessor of h that shares an information set

with h , which is well-defined in a tree as the predecessors of h are well-ordered. If $\mathcal{E}(h) \prec h$, let $\alpha(h)$ denote the unique action at $\mathcal{E}(h)$ that yields a history which precedes h .

DEFINITION: Transformation \mathcal{T}_i .

- For each $i \in N$, define $\mathcal{T}_i(G)$ by $G' = \langle H', \preceq', A', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g' \rangle$:
 - (i) $H'_i \equiv \{h \in H_i \mid \mathcal{E}(h) \prec h\}$,
 - (ii) $E \equiv \cup_{h \in H'_i} \{h' \in \sigma(h) \mid \mathcal{A}(h') \neq \alpha(h)\}$,
 - (iii) $H^* \equiv H \setminus \{h \in H \mid \text{there is } e \in E \text{ such that } e \preceq h\}$,
 - (iv) $H^{**} \equiv \{h \in H_i \mid |\sigma(h)| = 1\}$,
 - (v) $H' \equiv H^* \setminus H^{**}$, and
 - (vi) everything else is the associated restriction.
- Given \mathbb{S} for G , let $\mathcal{T}_i(\mathbb{S})$ be the associated restriction for $\mathcal{T}_i(G)$.

LEMMA 2.4: If $G \in \Gamma$ \mathbf{OSP}^S -implements f through \mathbb{S} , then for each $i \in N$, $\mathcal{T}_i(G)$ \mathbf{OSP}^S -implements f through $\mathcal{T}_i(\mathbb{S})$.

Again, we omit the proof as the notation obscures the logic. Suppose agent i has some history h that is strictly preceded by other histories in its information set, the earliest of which is $\mathcal{E}(h)$. If i has any pure strategies that can reach h ,¹² then all prescribe the action $\alpha(h)$ at $\mathcal{E}(h)$, and thus also prescribe $\alpha(h)$ at h . Any histories that follow from a choice other than $\alpha(h)$ at h , then, are impossible, and are pruned in the transformation. After the pruning, h only has one remaining immediate successor, and at this point is squished. Thus the transformation is again a composition of a squishing and a pruning, and in the transformed game form, i is not absent-minded. The rule's recommendation is still reached at each type profile, as only impossible plays have been removed. Moreover, i evaluates each strategy at fewer information sets, and as in the previous transformation, for each agent there are fewer possibilities at each information set; altogether, \mathbf{OSP}^S -implementation is again preserved. The composition of the \mathcal{T}_i transformations across N yields a game form with no absent-mindedness where each history has at least two immediate successors.

Let $\Gamma^* \subseteq \Gamma$ be the subclass of game forms without absent-mindedness. At this point, we remove imperfect information entirely, a possibility first observed by [Ashlagi and Gonczarowski \(2016\)](#):

DEFINITION: Transformation \mathcal{T}_* .

- For each $i \in N$, define $\mathcal{T}_i(G)$ by $G' = \langle H, \preceq, A, \mathcal{A}, P, (\mathcal{I}'_i)_{i \in N}, g \rangle$:
 - (i) for each $i \in N$, $\mathcal{I}'_i \equiv \{\{h\} \mid h \in H_i\}$, and
 - (ii) everything else is the same.
- Given \mathbb{S} for G , let $\mathcal{T}_*(\mathbb{S})$ be the associated restriction for $\mathcal{T}_*(G)$.

¹²It could, of course, be the case that i cannot reach one of his histories h with any pure strategy. For example, suppose all of i 's histories share an information set; at each history, the available actions are Left, Center, and Right; and reaching h requires a Left and then a Right.

LEMMA 2.5 (ASHLAGI AND GONCZAROWSKI, 2016): If $G \in \Gamma^*$ \mathbf{OSP}^S -implements f through \mathbb{S} , then $\mathcal{T}_*(G)$ \mathbf{OSP}^S -implements f through $\mathcal{T}_*(\mathbb{S})$.

We omit this proof for the same reason we omitted the last two. Because there is no absent-mindedness, each history is a possibility compatible with some pure strategy profile. When information sets are split into singletons, then, each agent considers fewer possibilities at each information set; as with the previous transformations, this implies that \mathbf{OSP}^S -implementation is preserved.

We remark that Lemma 2.5 would not hold if the original game form included absent-mindedness, as in this case the splitting of information sets could create new possibilities for agents to consider:

EXAMPLE 2: There is one agent with one possible type, and there are two outcomes, x and x' . The agent prefers x , but the rule f specifies that x' should be selected. The extensive game form G is borrowed from the absent-minded driver game (Piccione and Rubinstein, 1997): the agent first selects whether to stay or exit; if he exits, the game ends; if he stays, then he selects a second time whether to stay or exit, and in either case the game ends. The agent is absent-minded: both histories share an information set. The outcome function assigns x only to stay-exit, and assigns x' otherwise. Though G \mathbf{OSP}^S -implements f through the pure exit strategy, $\mathcal{T}_*(G)$ does not \mathbf{OSP}^S -implement f .

Thus the order in which these transformations are applied is important. That said, if they are applied in the order they are presented, the resulting game form is one of perfect information. At this point, we can apply a final transformation due to Li (2016) for game forms with perfect recall:

DEFINITION: Transformation $\mathcal{T}_{\mathbb{S}}$.

- For each \mathbb{S} , define $\mathcal{T}_{\mathbb{S}}(G)$ by $G' = \langle H', \lesssim', A', \mathcal{A}', P', (\mathcal{I}'_i)_{i \in N}, g' \rangle$:
 - (i) $H' \equiv \{h \in H \mid \text{there is } \theta \in \Theta \text{ such that } h \in \pi^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N})\}$, and
 - (ii) everything else is the associated restriction.
- Given \mathbb{S} for G , let $\mathcal{T}_{\mathbb{S}}(\mathbb{S})$ be the associated restriction for $\mathcal{T}_{\mathbb{S}}(G)$.

THEOREM L (Li, 2016): If perfect recall G \mathbf{OSP}^S -implements f through \mathbb{S} , then $\mathcal{T}_{\mathbb{S}}(G)$ \mathbf{OSP}^S -implements f through $\mathcal{T}_{\mathbb{S}}(\mathbb{S})$.

This last result allows us to use \mathbb{S} to prune off unused parts of the game tree, and therefore to label each action with the collection of types who take it. Altogether, we have established:

PROPOSITION 2: If a deterministic rule has an \mathbf{OSP}^S -implementation, then it has a round table implementation, which is an \mathbf{OSP}^C -implementation.

Appendix C

In this appendix, we prove Theorem 1:

THEOREM 1: If a stochastic rule has a random \mathbf{OSP}^S -implementation, then it has a randomized round table implementation, which is a random \mathbf{OSP}^C -implementation.

PROOF: Suppose (G, \mathcal{A}, μ) random \mathbf{OSP}^S -implements f through \mathbb{S} .

Let $\omega_{H_0} \in \Omega_{H_0}$. As argued for [Lemma 2.3](#), fixing this realization of the administrator's states, we may first remove all successors of any node that the administrator passes over, then squish the administrator's nodes, resulting in an admissible game tree and preserving random \mathbf{OSP}^S -implementation. Unlike [Lemma 2.3](#), however, the transformed game random \mathbf{OSP}^S -implements a *different* rule: we previously had implementation of a stochastic rule, and now have implementation of a deterministic rule. Relabel each remaining history h as $h^{\omega_{H_0}}$, let $(\mathbb{S}_i^{\omega_{H_0}})_{i \in N}$ denote the associated restriction of \mathbb{S} , and let $G^{\omega_{H_0}}$ denote the transformed game form, which has no administrator moves and \mathbf{OSP}^S -implements the deterministic rule $f^{\omega_{H_0}}$ given by:

$$f^{\omega_{H_0}}(\theta) = x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})).$$

Define $G^+ \in \Gamma$ such that (i) at the initial history, the administrator selects from Ω_{H_0} , (ii) each initial action ω_{H_0} is followed by the game form $G^{\omega_{H_0}}$, which has no administrator moves, and (iii) the administrator's strategy B_0^+ is the identity. For each $i \in N$, each $\theta_i \in \Theta_i$ and each of i 's histories $h^{\omega_{H_0}}$, define $\mathbb{S}_i^+(\theta_i)(h^{\omega_{H_0}}) \equiv \mathbb{S}_i^{\omega_{H_0}}(\theta_i)(h^{\omega_{H_0}})$. Clearly, for each $\theta \in \Theta$, $(\mathbb{S}_i^+(\theta_i)) \in \mathbf{OSP}^S(G^+, \theta)$, since at each information set, each agent compares the same possibilities that he did at one of the game forms $G^{\omega_{H_0}}$.

Consider $(G^+, \mathcal{A}, \mu) \in \Gamma^{MP}$. By construction, for each $\theta \in \Theta$ and each $X' \in \mathcal{X}$, the administrator states that lead to X' -outcomes in G^+ are precisely those that label actions whose ensuing game forms lead to X' -outcomes in G^+ , which are precisely those that lead to X' -outcomes in G :

$$\begin{aligned} & \{\omega_{H_0} \in \Omega_{H_0} | x^{G^+}(h_\emptyset^+, (\mathbb{S}_i^+(\theta_i))_{i \in N}, B_0^+(\omega_{H_0})) \in X'\} \\ &= \{\omega_{H_0} \in \Omega_{H_0} | x^{G^{\omega_{H_0}}}(h_\emptyset^{\omega_{H_0}}, (\mathbb{S}_i^{\omega_{H_0}}(\theta_i))_{i \in N}) \in X'\} \\ &= \{\omega_{H_0} \in \Omega_{H_0} | x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}, B_0(\omega_{H_0})) \in X'\}. \end{aligned}$$

Because μ assigns the same number to the former in (G^+, \mathcal{A}, μ) as to the latter in (G, \mathcal{A}, μ) , thus

$$\mu(\{\omega_{H_0} \in \Omega_{H_0} | x^{G^+}(h_\emptyset^+, (\mathbb{S}_i^+(\theta_i))_{i \in N}, B_0^+(\omega_{H_0})) \in X'\}) = f(\theta)(X'),$$

so (G^+, \mathcal{A}, μ) random \mathbf{OSP}^S -implements f .

It remains to transform (G^+, \mathcal{A}, μ) into a randomized round table mechanism (G^*, \mathcal{A}, μ) while preserving random \mathbf{OSP}^S -implementation of f . To do so, simply use [Proposition 2](#) to replace each game form $G^{\omega_{H_0}}$ with a round table mechanism that \mathbf{OSP}^S -implements the same deterministic rule $f^{\omega_{H_0}}$, and let $(\mathbb{S}_i^*)_{i \in N}$ gather the associated strategies. By construction,

$$\begin{aligned} & \{\omega_{H_0} \in \Omega_{H_0} | x^{G^*}(h_\emptyset^*, (\mathbb{S}_i^*(\theta_i))_{i \in N}, B_0^*(\omega_{H_0})) \in X'\} \\ &= \{\omega_{H_0} \in \Omega_{H_0} | x^{G^+}(h_\emptyset^+, (\mathbb{S}_i^+(\theta_i))_{i \in N}, B_0^+(\omega_{H_0})) \in X'\}, \end{aligned}$$

and μ assigns the same number to the former in (G^*, \mathcal{A}, μ) as to the latter in (G^+, \mathcal{A}, μ) , so (G^*, \mathcal{A}, μ) random \mathbf{OSP}^S -implements f , as desired. That (G^*, \mathcal{A}, μ) is a randomized round table mechanism is obvious. ■

Appendix D

In this appendix, we prove [Theorem 2](#):

THEOREM 2: Random serial dictatorship has no round table implementation.

PROOF: Let f denote random serial dictatorship. Assume, by way of contradiction, that there is round table mechanism G that \mathbf{OSP}^S -implements f .

Let us say that a history is *nondegenerate* if it has more than one action.¹³ As G \mathbf{OSP}^S -implements f , there must be a sharp history h . As the game tree is a *tree*, the predecessors of h are well-ordered, and thus h has an earliest sharp predecessor h^* . As G is a round table mechanism, h^* is controlled by some agent $i \in N$ and the actions at h^* are labeled by sets of types; as h^* has no sharp predecessors, these actions partition Θ_i .

Define θ_i^* by:

$$\theta_i^*(a) \equiv \begin{cases} 9, & a = A, \\ 6, & a = B, \\ 0, & a = C. \end{cases}$$

Let Θ_* be the action at h^* that includes θ_i^* , and let Θ' be a distinct action at h^* .

Let us evaluate $\mathbb{S}_i(\theta_i^*)$ at h^* , supposing that i has type θ_i^* . As h^* has no sharp predecessors and G \mathbf{OSP}^S -implements f , thus the worst-case scenario from continuing with $\mathbb{S}_i(\theta_i^*)$ is an outcome where i receives a uniform lottery over the three objects. Thus Θ' cannot include a type where A is top-ranked, else by similar logic the best-case scenario from deviating to Θ' would be an outcome where i receives A with probability 1, contradicting that $\mathbb{S}_i(\theta_i^*)$ is obviously dominant for type θ_i^* . Moreover, Θ' cannot include a type where B is top-ranked, else by deviating it would be possible for i to receive B with probability 1, contradicting that $\mathbb{S}_i(\theta_i^*)$ is obviously dominant for type θ_i^* . As Θ' was arbitrary, thus all actions at h^* except for Θ_* only include types where C is top-ranked.

Define θ_i^{**} by:

$$\theta_i^{**}(a) \equiv \begin{cases} 9, & a = A, \\ 6, & a = C, \\ 0, & a = B. \end{cases}$$

Let us evaluate $\mathbb{S}_i(\theta_i^{**})$ at h^* , supposing that i has type θ_i^{**} . Since θ_i^{**} does not top-rank C , thus $\theta_i^{**} \in \Theta_*$. Using the logic from the previous paragraph, the worst-case scenario from continuing with $\mathbb{S}_i(\theta_i^{**})$ is an outcome where i receives a uniform lottery over the three objects, and by deviating to Θ' (where necessarily all types top-rank C) it would be possible for i to receive C with probability 1. But then $\mathbb{S}_i(\theta_i^{**})$ is not obviously dominant for type θ_i^{**} , contradicting that G \mathbf{OSP}^S -implements f . ■

Appendix E

In this appendix, we prove [Theorem 3](#).

We begin with two similar observations. First, if a deterministic rule has an \mathbf{OSP}^M -implementation, then it has one in pure strategies:

¹³[Bade and Gonczarowski \(2017\)](#) observe that attention can be restricted to game trees for which all histories are nondegenerate.

LEMMA 3.1: If $G \in \Gamma$ $\mathbf{OSP}^{\mathcal{M}}$ -implements f through \mathbb{M} , then for each $\omega_N \in \Omega_N$, G $\mathbf{OSP}^{\mathcal{S}}$ -implements f through \mathbb{S} , where for each $i \in N$ and each $\theta_i \in \Theta_i$, $\mathbb{S}_i(\theta_i) = \mathbb{M}_i(\theta_i)(\omega_i)$.

PROOF: Suppose G $\mathbf{OSP}^{\mathcal{M}}$ -implements f through \mathbb{M} , and let $\omega_N \in \Omega_N$. For each $i \in N$ and each $\theta_i \in \Theta_i$, define $\mathbb{S}_i(\theta_i) \equiv \mathbb{M}(\theta_i)(\omega_i)$.

Let $\theta \in \Theta$ and let $i \in N$. Define $M_i \equiv \mathbb{M}_i(\theta_i)$ and define $S_i \equiv \mathbb{S}_i(\theta_i)$. We claim S_i is obviously \mathcal{S} -dominant. Indeed, let $I_i \in \mathcal{I}_i(S_i)$. Since $S_i = M_i(\omega_i)$, thus $I_i \in \mathcal{I}_i(M_i(\omega_i))$. Since M_i is obviously dominant, for each pair $h, h' \in I_i$, each $M'_i \in \mathcal{M}_i$ such that $M_i(\omega_i)(I_i) \neq M'_i(\omega_i)(I_i)$, each pair $M_{-i}, M'_{-i} \in \mathcal{M}_{-i}$, each pair $\omega_{-i}, \omega'_{-i} \in \Omega_{-i}$, and each pair $\omega_{H_0}, \omega'_{H_0} \in \Omega_{H_0}$,

$$\begin{aligned} & x^G(h, M_i(\omega_i), (M_j(\omega_j))_{j \in N \setminus \{i\}}, B_0(\omega_{H_0})) \\ & \quad R_i \\ & x^G(h', M'_i(\omega_i), (M'_j(\omega'_j))_{j \in N \setminus \{i\}}, B_0(\omega'_{H_0})). \end{aligned}$$

Moreover, $S_i = M_i(\omega_i)$, and each pure strategy is a realization of some mixed strategy. Thus for each pair $h, h' \in I_i$, each $S'_i \in \mathcal{S}_i$ such that $S_i(I_i) \neq S'_i(I_i)$, each pair $S_{-i}, S'_{-i} \in \mathcal{S}_{-i}$, and each pair $\omega_{H_0}, \omega'_{H_0} \in \Omega_{H_0}$,

$$\begin{aligned} & x^G(h, S_i, S_{-i}, B_0(\omega_{H_0})) \\ & \quad R_i \\ & x^G(h', S'_i, S'_{-i}, B_0(\omega'_{H_0})), \end{aligned}$$

so S_i is obviously dominant. Since $i \in N$ was arbitrary, thus $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^{\mathcal{S}}(G, \theta)$. Since $\theta \in \Theta$ was arbitrary, thus G $\mathbf{OSP}^{\mathcal{S}}$ -implements f through \mathbb{S} . Since $(\omega_i) \in \times \Omega_i$ was arbitrary, we are done. ■

Second, if a deterministic rule has an $\mathbf{OSP}^{\mathcal{B}}$ -implementation, then it has one in pure strategies:

LEMMA 3.2: If $G \in \Gamma$ $\mathbf{OSP}^{\mathcal{B}}$ -implements f through \mathbb{B} , then for each $\omega_{H_N} \in \Omega_{H_N}$ such that for each pair h and h' that share an information set, $\omega_h = \omega_{h'}$, G $\mathbf{OSP}^{\mathcal{S}}$ -implements f through \mathbb{S} , where for each $i \in N$ and each $\theta_i \in \Theta_i$, $\mathbb{S}_i(\theta_i) = \mathbb{B}_i(\theta_i)(\omega_{H_i})$.

PROOF: Suppose G $\mathbf{OSP}^{\mathcal{B}}$ -implements f through \mathbb{B} , and let $\omega_{H_N} \in \times \Omega_{H_N}$ such that for each pair h and h' that share an information set, $\omega_h = \omega_{h'}$. For each $i \in N$ and each $\theta_i \in \Theta_i$, define $\mathbb{S}_i(\theta_i) \equiv \mathbb{B}(\theta_i)(\omega_{H_i})$.

First, we claim that \mathbb{S} is a Harsanyi pure strategy profile. Indeed, let $i \in N$, let $\theta_i \in \Theta_i$, let $I_i \in \mathcal{I}_i$, and let $h, h' \in I_i$. Since $\mathbb{B}_i(\theta_i)$ is a behavioral strategy, it includes the same mapping at h and h' ; since by hypothesis $\omega_h = \omega_{h'}$, thus $\mathbb{B}_i(\theta_i)(\omega_{H_i})(h_i) = \mathbb{B}_i(\theta_i)(\omega_{H_i})(h'_i)$. Since $I_i \in \mathcal{I}_i$ and $h, h' \in I_i$ were arbitrary, thus the choice function $\mathbb{B}_i(\theta_i)(\omega_{H_i})$ is a pure strategy. Since $\theta_i \in \Theta_i$ was arbitrary, thus \mathbb{S}_i is a Harsanyi pure strategy. Since $i \in N$ was arbitrary, thus \mathbb{S} is a Harsanyi pure strategy profile, as desired.

Let $\theta \in \Theta$ and let $i \in N$. Define $B_i \equiv \mathbb{B}_i(\theta_i)$ and define $S_i \equiv \mathbb{S}_i(\theta_i)$. We claim S_i is obviously \mathcal{S} -dominant. Indeed, let $I_i \in \mathcal{I}_i(S_i)$. By construction, $\mathcal{I}_i(S_i) \subseteq \mathcal{I}_i(B_i)$, so $I_i \in \mathcal{I}_i(B_i)$. Since B_i is obviously dominant, for each pair $h, h' \in I_i$, each $\omega_x \in \Omega_h = \Omega_{h'}$, each $B'_i \in \mathcal{B}_i$ such that $B_i(\omega_x) \neq B'_i(\omega_x)$, each $\omega_{H_i \setminus \{h\}} \in \Omega_{H_i \setminus \{h\}}$, each $\omega'_{H_i \setminus \{h\}} \in \Omega_{H_i \setminus \{h\}}$, each pair $B_{-i}, B'_{-i} \in \mathcal{B}_{-i}$, each pair $(\omega_{H_j}^*)_{j \in N \setminus \{i\}}, (\omega'_{H_j}) \in \times_{N \setminus \{i\}} \Omega_{H_j}$, and each pair $\omega_{H_0}, \omega'_{H_0} \in \Omega_{H_0}$,

$$x^G(h, B_i(\omega_x, \omega_{H_i \setminus \{h\}}^*), (B_j(\omega_{H_j}))_{j \in N \setminus \{i\}}, B_0(\omega_{H_0})) \\ R_i \\ x^G(h', B'_i(\omega_x, \omega'_{H_i \setminus \{h'\}}), (B'_j(\omega'_{H_j}))_{j \in N \setminus \{i\}}, B_0(\omega'_{H_0})).$$

In particular, this holds for $(\omega_x, \omega_{H_i \setminus \{h\}}^*) = \omega_{H_i}$. Moreover, each pure strategy is a realization of some behavioral strategy. Thus for each pair $h, h' \in I_i$, each $S'_i \in \mathcal{S}_i$ such that $S_i(I_i) \neq S'_i(I_i)$, each pair $S_{-i}, S'_{-i} \in \mathcal{S}_{-i}$, and each pair $\omega_{H_0}, \omega'_{H_0} \in \Omega_{H_0}$,

$$x^G(h, S_i, S_{-i}, B_0(\omega_{H_0})) \\ R_i \\ x^G(h', S'_i, S'_{-i}, B_0(\omega'_{H_0})),$$

so S_i is obviously dominant. Since $i \in N$ was arbitrary, thus $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^{\mathcal{S}}(G, \theta)$. Since $\theta \in \Theta$ was arbitrary, thus G $\mathbf{OSP}^{\mathcal{S}}$ -implements f through \mathbb{S} . Since $(\omega_{H_N}) \in \times \Omega_{H_N}$ satisfying the given restriction was arbitrary, we are done. ■

By [Lemma 3.1](#), [Lemma 3.2](#), and [Proposition 2](#), if a rule has any kind of obviously strategy-proof implementation, then it has a round table implementation. Altogether, we have established:

THEOREM 3: If a deterministic rule f is either

- (i) $\mathbf{OSP}^{\mathcal{S}}$ -implementable,
- (ii) $\mathbf{OSP}^{\mathcal{M}}$ -implementable, or
- (iii) $\mathbf{OSP}^{\mathcal{B}}$ -implementable,

then it has a round table implementation, which is an $\mathbf{OSP}^{\mathcal{C}}$ -implementation.

Appendix F

In this appendix, we prove [Theorem 4](#):

THEOREM 4: Suppose there is a finite set of outcomes and each agent has a von Neumann-Morgenstern type compatible with a strict ranking over outcomes. If a stochastic rule f has a random $\mathbf{OSP}^{\mathcal{S}}$ -implementation, then f is ordinal and has a randomized ordinal round table implementation.

PROOF: Suppose f has a random $\mathbf{OSP}^{\mathcal{S}}$ -implementation. Then by [Theorem 1](#), it has a random round table implementation.

Let $i \in N$, and let $\theta_i, \theta'_i \in \Theta_i$ be a pair of von Neumann-Morgenstern utility functions that are compatible with the same ordinal preferences over outcomes. Let H_{θ_i, θ'_i} be the collection of histories where i has one action that includes θ_i and another action that includes θ'_i . At each $h \in H_{\theta_i, \theta'_i}$, let $X_{\theta_i}^h$ be the collection of outcomes that could occur when the strategy of reporting θ_i is followed from h , and let $X_{\theta'_i}^h$ be the collection of outcomes that could occur when the strategy of reporting θ'_i is followed from h . As truthful reporting of θ_i and truthful reporting of θ'_i are both obviously dominant for i , and as both of these types share a strict ranking of outcomes, it follows that $X_{\theta_i}^h$ and $X_{\theta'_i}^h$ are the same singleton.

Let $\theta_{-i} \in \Theta_{-i}$. By the above observation, the distribution over outcomes generated by administrator randomization for reports of (θ_i, θ_{-i}) is the same as the distribution

generated for (θ'_i, θ_{-i}) . As we have a random round table implementation of f , thus $f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i})$; it follows that f is ordinal.

Let $[\theta_i] \subseteq \Theta_i$ be the collection of types for i that are compatible with the same ranking of outcomes as θ_i . For each action that includes θ_i , relabel the action to be the union of its previous types with the types in $[\theta_i]$. For each action that does not include θ_i , relabel the action to be its previous types except for those in $[\theta_i]$. Finally, prune off all histories that follow from an action whose label is now empty. For each $\theta_i^* \in \Theta_i$ and each $\theta_{-i} \in \Theta_{-i}$, the reports $(\theta_i^*, \theta_{-i})$ determine the same outcome for each round table mechanism selected by the administrator before and after this operation; thus the distribution generated by administrator randomization for reports $(\theta_i^*, \theta_{-i})$ is the same before and after this operation, so we have a new round table implementation of f .

This relabeling and pruning can be done for each of the finitely many rankings for i , and then all of these stages of relabeling and pruning can be done for each of the finitely many agents. Finally, relabel each action by its associated rankings of its types; the result is clearly a random ordinal round table implementation of f . ■

Appendix G

In this appendix, we prove [Theorem 5](#):

THEOREM 5: A deterministic rule f has an \mathbf{OSP}^C -implementation if and only if it has a \mathbf{SP} -implementation through a round table mechanism.

PROOF: We prove the two implications in sequence. Note that because for each round table mechanism $G \in \Gamma^{RT}$, the administrator does not play ($H_0 = \emptyset$), we will suppress notation involving him.

[\Rightarrow] Assume f has an \mathbf{OSP}^C -implementation. By [Proposition 2](#), there are $G \in \Gamma^{RT}$ and an associated Harsanyi pure strategy profile \mathbb{S} such that G \mathbf{OSP}^C -implements f through \mathbb{S} . For each $\theta \in \Theta$, since

- by definition, $(\mathbb{S}_i(\theta_i)) \in \mathbf{OSP}^C(G, \theta)$;
- because $G \in \Gamma^{RT}$, $\mathbf{OSP}^C(G, \theta) = \mathbf{OSP}^S(G, \theta)$; and
- by definition, $\mathbf{OSP}^S(G, \theta) \subseteq \mathbf{SP}(G, \theta)$;

thus altogether $(\mathbb{S}_i(\theta_i)) \in \mathbf{SP}(G, \theta)$. Since for each $\theta \in \Theta$, $x^G(h_\emptyset, (\mathbb{S}_i(\theta_i))_{i \in N}) = f(\theta)$, thus G \mathbf{SP} -implements f through \mathbb{S} , as desired.

[\Leftarrow] Assume there are $G \in \Gamma^{RT}$ and an associated Harsanyi pure strategy profile \mathbb{S} such that G \mathbf{SP} -implements f through \mathbb{S} . We claim that G \mathbf{OSP}^C -implements f through \mathbb{S} . Indeed, let $\theta \in \Theta$, let $i \in N$, and define $S_i \equiv \mathbb{S}_i(\theta_i)$.

We claim S_i is obviously dominant. Indeed, let $I_i \in \mathcal{I}_i(S_i)$, let $h, h' \in \mathcal{I}_i$, let $C'_i \in \mathcal{C}_i$ such that $S_i(I_i) \neq C'_i(I_i)$, and let $C_{-i}, C'_{-i} \in \mathcal{C}_{-i}$. (As $G \in \Gamma^{RT}$, the administrator does not play; $\mathcal{C}_0 = \emptyset$.)

Since $G \in \Gamma^{RT}$, $h = h'$. Define h_1 to be the unique next history after h determined by the action $S_i(h)$, and define h_2 to be the unique next history after h determined by the action $C'_i(h)$. Partition H_{-i} into H_0, H_1, H_2 , and H_* , where:

- $H_0 \equiv \{h' \in H_{-i} | h' \prec h\}$,
- $H_1 \equiv \{h' \in H_{-i} | h_1 \preceq h'\}$,
- $H_2 \equiv \{h' \in H_{-i} | h_2 \preceq h'\}$, and
- $H_* \equiv H_{-i} \setminus (H_1 \cup H_2)$.

We use these to define $S_{-i} \in S_{-i}$. Since $G \in \Gamma^{RT}$, we can do so history-by-history:

- for each $h' \in H_0$, define $j \equiv P(h')$, and define $S_j(h')$ to be the action that determines the unique next history in the (well-ordered) chain $\{h'' \in H | h'' \preceq h\}$;
- for each $h' \in H_1$, define $j \equiv P(h')$, and define $S_j(h') \equiv C_j(h')$;
- for each $h' \in H_2$, define $j \equiv P(h')$, and define $S_j(h') \equiv C'_j(h')$; and
- for each $h' \in H_*$, define $j \equiv P(h')$, and let $S_j(h') \in A(h')$ be arbitrary.

Define $S'_i \in \mathcal{S}_i$ such that

- for each $h' \in H_i$ such that $h' \prec h$, $S'_i(h') = S_i(h)$;
- for each $h' \in H_i$ such that $h \preceq h'$, $S'_i(h') = C'_i(h')$; and
- for all other h' , $S'_i(h') \in A(h')$ is arbitrary.

Since S_i is dominant, we have

$$\begin{aligned} & x^G(h_\emptyset, S_i, S_{-i}) \\ & \quad R_i \\ & x^G(h_\emptyset, S'_i, S_{-i}). \end{aligned}$$

Since $\{h\} = I_i \in \mathcal{I}_i(S_i)$, by construction of S_{-i} , $h \in \pi^G(h_\emptyset, S_i, S_{-i})$. Since S_i and S'_i agree on histories that precede h , thus $h \in \pi^G(h_\emptyset, S'_i, S_{-i})$. Therefore, by construction of S_{-i} ,

- (i) $x^G(h_\emptyset, S_i, S_{-i}) = x^G(h, S_i, C_{-i})$, and
- (ii) $x^G(h_\emptyset, S'_i, S_{-i}) = x^G(h', C'_i, C'_{-i})$,

so S_i is obviously \mathcal{C} -dominant. Since $i \in N$ was arbitrary, $(S_i(\theta_i)) \in \mathbf{OSP}^{\mathcal{C}}(G, \theta)$. Since $\theta \in \Theta$ was arbitrary, thus for each $\theta \in \Theta$, $(S_i(\theta_i)) \in \mathbf{OSP}^{\mathcal{C}}(G, \theta)$. Since for each $\theta \in \Theta$, $x^G(h_\emptyset, (S_i(\theta_i))) = f(\theta)$, thus G $\mathbf{OSP}^{\mathcal{C}}$ -implements f through \mathbb{S} , as desired. ■

References

- ALÓS-FERRER, C., KERN, J., AND RITZBERGER, K. (2011). “Comment on ‘Trees and extensive forms.’” *Journal of Economic Theory* 146, 2165-2168.
- ALÓS-FERRER, C. AND RITZBERGER, K. (2005a). “Some remarks on pseudotrees.” *Order* 22, 1-9.
- ALÓS-FERRER, C. AND RITZBERGER, K. (2005b). “Trees and decisions.” *Economic Theory* 25, 763-798.
- ALÓS-FERRER, C. AND RITZBERGER, K. (2008). “Trees and extensive forms.” *Journal of Economic Theory* 143, 216-250.

- ALÓS-FERRER, C. AND RITZBERGER, K. (2016). *The Theory of Extensive Form Games*. Berlin, Germany: Springer-Verlag Berlin Heidelberg.
- ARRIBILLAGA, R., MASSÓ, J., AND NEME, A. (2017). “Not All Majority-based Social Choice Functions Are Obviously Strategy-proof.” Working paper.
- ASHLAGI, I. AND GONCZAROWSKI, Y. (2016). “Stable matching mechanisms are not obviously strategy-proof.” Working paper.
- AUMANN, R. (1964). “Mixed and Behavior Strategies in Infinite Extensive Games.” In: *Advances in Game Theory*, Annals of Mathematics Studies 52. Editors: Dresher, M., Shapley, L., and Tucker, A. Princeton, New Jersey: Princeton University Press.
- AUMANN, R. (1987). “Correlated Equilibrium as an Expression of Bayesian Rationality.” *Econometrica* 55, 1-18.
- BADE, S. AND GONCZAROWSKI, Y. (2017). “Gibbard-Satterthwaite Success Stories and Obvious Strategy-proofness.” Working paper.
- CARROLL, G. (2017). “On Mechanisms Eliciting Ordinal Preferences.” Forthcoming, *Theoretical Economics*.
- FRANK, A. (1947). *Het Achterhuis* [*The Diary*, in Dutch]. Amsterdam, the Netherlands: Contact Publishing. Translation: Frank, A. and Massotty, S. (1995). *The Diary of a Young Girl: The Definitive Edition*. Editors: Frank, O. and Pressler, M. New York, New York: Doubleday.
- GALE, D. AND SHAPLEY, L. (1962). “College Admissions and the Stability of Marriage.” *The American Mathematical Monthly* 69, 9-15.
- GIBBARD, A. (1973). “Manipulation of Voting Schemes: A General Result.” *Econometrica* 41, 587-601.
- HALMOS, P. (1950). *Measure Theory*. New York, New York: Van Nostrand.
- HARSANYI, J. (1967). “Games with incomplete information played by ‘Bayesian’ players, I-III: Part I, The Basic Model.” *Management Science* 14, 159-182.
- ISELL, J. (1957). “Finitary games.” In: *Contributions to the Theory of Games, Volume III*, Annals of Mathematics Studies 39. Editors: Dresher, M., Tucker, A., and Wolfe, P. Princeton, New Jersey: Princeton University Press.
- JECH, T. (1971). “Trees.” *The Journal of Symbolic Logic* 36, 1-14.
- KOPPELBERG, S. AND MONK, J. (1991). “Pseudo-trees and Boolean algebras.” *Order* 8, 359-374.
- KUHN, H. (1953). “Extensive games and the problem of information.” In: *Contributions to the Theory of Games, Volume II*, Annals of Mathematics Studies 28. Editors: Kuhn, H. and Tucker, A. Princeton, New Jersey: Princeton University Press.
- LI, S. (2016). “Obviously Strategy-Proof Mechanisms.” *American Economic Review* 107, 3257-3287.

- MACKEY, G. (1957). "Borel structure in groups and their duals." *Transactions of the American Mathematical Society* 85, 134-165.
- MILGROM, P. AND WEBER, R. (1982). "A Theory of Auctions and Competitive Bidding." *Econometrica* 50, 1089-1122.
- MYERSON, R. (1981). "Optimal Auction Design." *Mathematics of Operation Research* 6, 58-73.
- OSBORNE, M. AND RUBINSTEIN, A. (1994). *A Course in Game Theory*. Cambridge, Massachusetts: The MIT Press.
- PICCIONE, M. AND RUBINSTEIN, A. (1997). "On the Interpretation of Decision Problems with Imperfect Recall." *Games and Economic Behavior* 20, 3-24.
- PYCIA, M. AND TROYAN, P. (2017). "Obvious Dominance and Random Priority." Working paper.
- SHAPLEY, L. AND SCARF, H. (1974). "On cores and indivisibility." *Journal of Mathematical Economics* 1, 23-37.
- TROYAN, P. (2016). "Obviously Strategy-Proof Implementation of Top Trading Cycles." Working paper.
- VICKREY, W. (1961). "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance* 16, 8-37.