

Are Simple Mechanisms Optimal when Agents are Unsophisticated?*

Jiangtao Li [Ⓔ] Piotr Dworczak[†]

January 19, 2022

Abstract

We study the design of mechanisms involving agents that have limited strategic sophistication. We define a mechanism to be *simple* if—given the assumed level of strategic sophistication—agents can determine their optimal strategy. We examine whether it is *optimal* for the mechanism designer who faces strategically unsophisticated agents to offer a simple mechanism. We show that when the designer uses a mechanism that is not simple, while she loses the ability to predict play, she may nevertheless be better off no matter how agents resolve their strategic confusion.

KEYWORDS: Simple mechanisms, complex mechanisms, robust mechanism design, dominant-strategy mechanisms, obviously strategy-proof mechanisms, strongly obviously strategy-proof mechanisms

*We thank Mohammad Akbarpour, Tilman Börgers, Gabriel Carroll, Jeff Ely, Scott Duke Kominers, Shengwu Li, Paul Milgrom, Stephen Morris, Parag Pathak, Andrzej Skrzypacz, Rani Spiegler, Satoru Takahashi, Takuro Yamashita, and participants at various seminars and conferences for helpful discussions. An abstract of this paper appeared in the *Proceedings of the 22nd ACM Conference on Economics and Computation (EC'21)*. Li thanks the Department of Economics at the Northwestern University for hospitality and support during the academic year 2018-2019, when some of the work on this paper was undertaken.

[†]The “[Ⓔ]” symbol indicates that the authors’ names are in certified random order, as described by Ray [Ⓔ] Robson (2018). *Li*: School of Economics, Singapore Management University; jtli@smu.edu.sg. *Dworczak*: Department of Economics, Northwestern University; GRAPE Research Center; piotr.dworczak@northwestern.edu.

1 Introduction

It is widely accepted that “real-life” economic agents are not as rational as their counterparts in economic models. When agents have limited strategic sophistication, economists lose confidence in the performance of mechanisms that force participants to engage in complicated mental tasks. For example, achieving a Bayesian Nash equilibrium requires each agent to know the distribution of other agents’ private information and correctly forecast other agents’ strategies; this is why dominant-strategy (or strategy-proof or SP) mechanisms are generally perceived as being superior for practical purposes. Following the mounting evidence that even dominant strategies are difficult to identify for real-life agents, several recent papers introduced classes of mechanisms in which agents can determine their optimal strategies under even weaker assumptions about strategic sophistication. [Li \(2017\)](#) proposes the notion of obviously strategy-proof (OSP) mechanisms in which agents can determine their optimal strategy even if they cannot engage in contingent reasoning. [Pycia and Troyan \(2021\)](#) strengthen the notion of simplicity even further by relaxing the assumption that agents can predict their own future moves, and define (among other intermediate concepts) strongly obviously strategy-proof (SOSP) mechanisms.

For the purpose of this paper, we call a mechanism *simple* if, given the assumed level of strategic sophistication, agents can determine their optimal strategy in the mechanism. For example, if we are only comfortable assuming that agents avoid obviously dominated strategies ([Li \(2017\)](#)), then an OSP mechanism is simple because an obviously dominant strategy can be identified as the unique strategy (up to payoff equivalence) that is not obviously dominated for such agents. If the designer instead offers a mechanism that is not OSP, she can no longer predict how agents will behave. More generally, we call a mechanism *complex* if it creates *strategic confusion* for the agents, understood as the inability to determine their optimal strategy in the mechanism.

The key observation of this paper is that the inability of the designer to predict the outcome of a complex mechanism need not be a sufficient reason for the use of simple mechanisms. As long as the designer is ultimately concerned with maximizing her own payoff—which is typically assumed in mechanism design—complex mechanisms may be preferred by the designer to simple ones.

Since a complex mechanism, by definition, leads to more than one possible outcome, we need to specify what we mean by the designer preferring a complex mechanism to a

simple one. We take a robust approach, defining *strong dominance* to be a case in which the complex mechanism generates a strictly higher expected payoff for the designer, no matter how agents resolve their strategic confusion. That is, even if agents choose their strategies to minimize the designer’s expected payoff whenever they were confused, the complex mechanism would still be strictly preferred by the designer to the simple one. We illustrate the possibility of strong dominance using an example.

Example 1. Consider the problem of designing a trading platform for two traders with the goal of maximizing intermediation profits. Each trader can buy or (short) sell one unit of the asset. The platform cannot hold inventory (ex-post market-clearing is imposed). Trader *A*’s valuation for the asset is either 0 or $2/3$. Trader *B*’s valuation for the asset is either $1/3$ or 1. The designer of the trading platform believes individual types to be equally likely but correlated across traders: $\pi((0, 1/3)) = \pi((2/3, 1)) = 2/5$.¹

Suppose that the traders are unable to predict each others’ strategies, and cannot engage in contingent reasoning. Then, as formalized and shown by Li (2017), traders will not play an obviously dominated strategy. Agents’ behavior is pinned down if and only if the designer uses an OSP mechanism, which is thus the relevant notion of a simple mechanism under this assumption about agents’ rationality.²

The optimal OSP mechanism yields an expected profit of $1/5$ for the platform and is implemented by the following extensive-form game, which can be viewed as an ascending personal-clock auction:

1. Trader *A* is asked whether she would like to sell the asset at the price 0; if she says “yes,” then that trade is implemented; if she says “no,” then:
2. Trader *B* is asked whether she would like to sell the asset at the price $1/3$; if she says “yes,” then that trade is implemented; if she says “no,” then there is no trade.

¹Because each trader’s role as a buyer or a seller is endogenously determined, the binding incentive constraints cannot be pinned down ex-ante. The correlation of types (and the specific value of $2/5$) plays no role in our analysis, except for ensuring that the type profile $(0, 1)$ is relatively unlikely—leading to a particular structure of binding IC and IR constraints in the optimal simple mechanism. For related models, see Cramton et al. (1987), Lu and Robert (2001), Chen and Li (2018), and Loertscher and Marx (2020).

²A strategy is obviously dominant if, for any deviation, at any information set where the two strategies first diverge, the best outcome under the deviation is no better than the worst outcome under the dominant strategy. A mechanism is OSP if it has an equilibrium in obviously dominant strategies. See Section 2 for a formal definition.

In all cases in which trade takes place, the platform charges a fee of $1/3$ to the buyer, that is, the buyer pays the trading price plus $1/3$.³

It is obviously dominant for type 0 of trader A to accept the initial offer, and for type $2/3$ to reject it. It is obviously dominant for type $1/3$ of trader B to accept the final offer, and for type 1 to reject it. It follows that the platform's profit is $1/3$ except when the type profile is $(2/3, 1)$. Intuitively, the inefficient no-trade outcome at the type profile $(2/3, 1)$ is implemented so that type 0 has an obviously dominant strategy: If trader B were buying the asset from trader A conditional on the profile $(2/3, 1)$, then the best possible outcome for type 0 from rejecting the initial offer would yield a strictly positive payoff, while her equilibrium strategy yields a payoff of 0.

Consider now an alternative mechanism, which can be viewed as a descending personal-clock auction:

1. Trader B is asked whether she would like to buy the asset at the price 1; if she says "yes," then that trade is implemented; if she says "no," then:
2. Trader A is asked whether she would like to buy the asset at the price $2/3$; if she says "yes," then that trade is implemented; if she says "no," then:
3. Trader B is asked whether she would like to buy the asset at the price $1/3$; if she says "yes," then that trade is implemented; if she says "no," then there is no trade.

Conditional on trade, the platform charges a fee of $1/3$ to the seller.

Both types of trader A and type $1/3$ of trader B have obviously dominant strategies; for example, it is obviously dominant for type $1/3$ of trader B to reject the first offer and to accept the second offer. But the mechanism is not OSP: Type 1 of trader B is confused between accepting the initial offer (which gives her 0) and rejecting it while accepting the second offer (which gives her $-2/3$ or $2/3$, depending on the behavior of trader A). The key observation is that, regardless of how trader B resolves this confusion, trade always happens. Thus, the platform achieves a profit of $1/3$ ex post, and hence also in expectation. Finally, for each type, non-participation is obviously dominated.⁴ Thus, as long as traders do not play obviously dominated strategies, by adopting a complex mechanism, the platform is guaranteed to achieve a strictly higher revenue than in the optimal OSP mechanism. \square

³We show the optimality of this mechanism in Appendix A.1.

⁴The platform could ensure that non-participation is obviously *strictly* dominated by adjusting the fee and the prices by an arbitrarily small $\epsilon > 0$.

While Example 1 focused on the solution concept of OSP, in the remainder of the paper we study a range of solution concepts (SP, OSP and SOSP), and derive three general results about strong dominance. First, we argue that simple mechanisms may be overly restrictive by requiring that no type be strategically confused. Sometimes, the designer may benefit by trading off better performance of simple mechanisms on a subset of the type space against creating confusion for the remaining types. We formalize this idea by defining a property called “accommodation of additional types (AAT).” We show that when the AAT property is violated, strong dominance of the best simple mechanism occurs for some objective function of the designer or some distribution of agents’ types. The AAT property fails in many classical social choice environments. As an application, we study a voting environment in which the designer attempts to maximize Rawlsian welfare of two agents with privately observed preferences over three alternatives. The best simple mechanism—serial dictatorship—is strongly dominated by a complex mechanism in which both agents influence the final choice of the alternative.

Second, we show that strong dominance is possible even when there is a single agent in the mechanism. In this case, the designer uses randomization to confuse the agent.⁵ Concretely, we construct an example in which the only way to implement a certain desirable random outcome is to create strategic confusion: After selecting a binary lottery from a menu, the agent is additionally asked to choose the side of the coin resulting in her preferred option—that choice makes the agent strategically confused. Remarkably, any mechanism that predetermines the consequences of the coin flip (any simple mechanism) leads to a strictly worse outcome for the designer. We extend the logic of the example to provide a sufficient condition for strong dominance in the single-agent environment.

Third, we identify environments in which simple mechanisms are *not* strongly dominated. We show that the optimal simple mechanism is not strongly dominated when the designer’s maximized objective function (under the chosen solution concept) is the same as the value of a relaxed problem in which the incentives constraints are only imposed along the edges of some directed tree in the type space. A notable instance of such a setting is the single-unit auction in the regular case. In this and other settings, the result establishes an *optimality foundation* for the use of simple mechanisms.

⁵Our results for a single agent rely on the assumption that the agent reasons about randomized outcomes in the same way as she reasons about choices of other players, see Remark 1.

The notion of strong dominance is quite demanding, as it requires that the superior complex mechanism generate a *strictly higher* expected payoff to the designer, regardless of how agents behave when they are confused. This property ensures robustness to how agents choose among undominated strategies; if agents turn out to be more sophisticated and rule out more strategies (than hypothesized by the designer), the performance guarantee of a complex mechanism can only improve. However, in some applications, the designer may be satisfied with a weaker property of a superior complex mechanism: It *never* yields a lower payoff than the best simple mechanism, and *sometimes* yields a strictly higher payoff. We call this property “weak dominance.”⁶

We argue that weak dominance of the best simple mechanism is a common phenomenon. We prove a general result showing how to construct a weakly-dominating complex mechanism from a given simple mechanism. Intuitively, the complex mechanism is created by augmenting the simple mechanism with an additional option that—if taken—benefits the designer. The additional option is made sufficiently attractive for the agent so that choosing it cannot be ruled out given the assumed level of sophistication. We show that this construction can be applied to a variety of settings, including revenue-maximization problems, efficiency in bilateral trade, and voting. The following example illustrates the logic behind the general result.

Example 2. Consider the problem of selling an item to one of N bidders with independent private values distributed according to a regular and symmetric distribution. Suppose that, as in Example 1, the only assumption about agents is that they do not play obviously dominated strategies. The best simple (OSP) mechanism is an ascending clock auction with a distribution-dependent reserve price. In the ascending clock auction, active bidders choose whether to exit as the clock price increases, and bidders who exit remain inactive thereafter. The auction stops when all but one bidder exit. The remaining bidder wins the object and pays the clock price.

The following modified mechanism, which we call the ascending clock auction with jump bidding, weakly dominates the best simple mechanism: Each bidder is allowed to speed up the clock by jump bidding, that is, to make a higher bid than the current

⁶It may be useful to draw an analogy to the notion of weak dominance between two strategies in game theory. The arguments for and against playing weakly dominated strategies carry over to selecting weakly dominated mechanisms. In particular, the designer could still perceive a weakly dominated mechanism to be optimal if she believes with certainty that agents choose strategies that are worst possible for her whenever they are confused, just as a weakly dominated strategy could be perceived to be optimal as it could be a best response to a degenerate belief about her opponent’s strategies.

clock price. This mechanism is not OSP because making a jump bid (to a bid b) is not obviously dominated for a bidder with value $v > b$ at the clock price $p < b$. Indeed, making a jump bid to b yields the best-case payoff of $v - b$ to the bidder, while following the default strategy (of exiting when the clock price reaches v) yields a payoff of 0 in the worst case.⁷ Agents are strategically confused as they now have multiple strategies that are not obviously dominated—the mechanism is complex.

Given the assumed strategic sophistication of the bidders, the designer cannot predict whether jump bidding will occur or not. Nonetheless, a revenue-maximizing auctioneer might prefer the ascending clock auction with jump bidding to the ascending clock auction: If none of the bidders jump bids, then the performance of the ascending clock auction with jump bidding is the same as that of the ascending clock auction; in the event that some bidder jump bids, the expected revenue of the ascending clock auction with jump bidding is strictly higher than that of the ascending clock auction. \square

We use a similar construction to show that the best simple mechanism is weakly dominated for a revenue-maximizing designer in a variety of settings. However, we also prove that, under some additional conditions, single-agent posted price mechanisms are *not* weakly dominated for the notions of SP and OSP (interestingly, this is not true under SOSP).

Overall, the paper proposes a systematic framework for thinking about the issues of strategic simplicity and complexity in settings with agents that are strategically unsophisticated. By offering results that both oppose and support the use of simple mechanisms, we emphasize that whether simplicity is desirable or not is not merely a function of agents' sophistication. Instead, each environment in question must be carefully analyzed before deciding whether a simple mechanism should be used.

1.1 Related literature

Simple mechanisms: This paper contributes to the design of mechanisms involving strategically unsophisticated agents. While [Li \(2017\)](#), [Börgers and Li \(2019\)](#), and [Pycia and Troyan \(2021\)](#) provide notions of simplicity and the characterizations of simple mechanisms according to these notions, our focus is on the tradeoff between simplicity and optimality, and we examine whether there is a foundation for the use of simple

⁷Of course, bidders would not jump bid if they could engage in contingent reasoning—jump bidding is weakly dominated but not obviously dominated.

mechanisms from an optimality perspective.⁸ Indeed, we show that in many cases, the designer might prefer a mechanism that is not simple.

Robust mechanisms design: Traditional models in mechanism design make strong assumptions about the detailed knowledge of the designer about the inputs to the mechanism design model. The literature of robust mechanism design seeks to relax these assumptions; see [Carroll \(2019\)](#) for a recent survey. While the leading interpretation of our exercise is that agents have limited strategic sophistication, an alternative interpretation—tying our work to the burgeoning literature of robust mechanism design—is that we relax the assumption about the designer’s knowledge of the strategic reasoning process of the agents (beyond some minimal rationality assumptions).

SP mechanisms: A number of papers have examined whether there is a foundation for the use of dominant-strategy mechanisms when the designer does not have any reliable information about the agents’ hierarchies of beliefs; see for example [Bergemann and Morris \(2005\)](#), [Chung and Ely \(2007\)](#), [Börgers \(2017\)](#), [Chen and Li \(2018\)](#), and [Yamashita and Zhu \(2020\)](#). In this approach, the agents can have arbitrary beliefs, while they are still assumed to play a Bayesian equilibrium. Our approach is different. While our analysis for the solution concept of SP also asks whether there is a foundation for the use of dominant-strategy mechanisms, we are not motivated by the robustness to the agents’ hierarchies of beliefs, but driven by simplicity concerns. We impose a minimal rationality assumption that agents do not play weakly dominated strategies; agents are not assumed to play a Bayesian equilibrium.⁹

Implementation in undominated strategies: Our analysis for the solution concept of SP is closely related to the strand of the mechanism design literature that studies implementation in undominated strategies; see, for example, [Börgers \(1991\)](#), [Jackson \(1992\)](#), [Börgers and Smith \(2012\)](#), [Carroll \(2014\)](#), [Yamashita \(2015a\)](#), [Mukherjee et al. \(2019\)](#), and [Mukherjee et al. \(2020\)](#). It is known from these papers that the optimal SP mechanism could be weakly or strongly dominated by complex mechanisms; we provide new examples and additional structural insights. For voting, [Börgers \(1991\)](#) shows that there are non-dictatorial procedures which ensure that collective decisions are Pareto efficient if all agents choose strategies that are not weakly dominated, and

⁸There are, of course, many dimensions of simplicity. Our paper and the papers cited here consider the strategic dimension.

⁹The concepts of structural and strategic uncertainty are related, as shown by [Yamashita \(2015b\)](#); the formal equivalence established by [Yamashita \(2015b\)](#) does not apply under our assumptions.

Mukherjee et al. (2019) show that the Pareto correspondence can be implemented in weakly undominated strategies by bounded mechanisms. Börgers and Smith (2012) show that, for bilateral trade and voting, there exist mechanisms that weakly dominate the optimal SP mechanism if all agents choose strategies that are not weakly dominated. In an independent paper, Mukherjee et al. (2020) provide sufficient conditions for a social choice correspondence to be implemented in weakly undominated strategies. They apply their results to several economic environments; their results are broadly consistent with and complementary to our Section 4 where we study the concept of weak dominance. Yamashita (2015a) develops a methodology of establishing an upper bound of the highest worst-case payoff for the designer, and applies the methodology to several settings, including the private-value auction and bilateral trade, and interdependent-value auction.¹⁰ Our results in Section 3.3 extend the insight of Yamashita (2015a) to a larger class of environments and the solution concepts of OSP and SOSP.

Simple versus complex mechanisms: There is a large computer-science literature—too vast to be surveyed here—that quantifies the worst-case loss from using certain simple classes of mechanisms relative to the optimal mechanism; for example, see Hartline and Roughgarden (2009). Pycia and Troyan (2021) ask whether restricting attention to simple mechanisms entails any loss by studying how the set of allocations that are simply implementable depends on agents’ degree of strategic sophistication. The fundamental difference to our work is that these papers compare the performance of mechanisms *across* solution concepts and the corresponding classes of mechanisms. We fix the assumption about agents’ rationality and compare the performance of simple mechanisms (in which agents know what to do) with the performance of complex mechanisms (in which agents are confused).

In terms of understanding complex mechanisms with cognitively limited agents, Jakobsen (2020) is similar in spirit to our paper. He studies a mechanism design problem involving a principal and a boundedly rational agent who has both imperfect memory and limited deductive (computational) ability. Thus, the agent’s comprehension of a game form—the mapping from strategy profiles to outcomes—is subject to complexity

¹⁰We are aware of two other papers that study questions different from ours but contain examples that could be used to show instances of strong dominance of an SP mechanism: Bergemann and Morris (2005)[Example 2] who study robust implementation of correspondences in settings in which the designer does not have information about agents’ belief and Carroll (2014) who proves a complexity result for undominated-strategy implementation. Also see Footnote 28 for the discussion of interdependent-value settings.

constraints. He shows that by expressing a mechanism as a complex contract, the principle could manipulate the agent into believing that truthful reporting is optimal. This paper differs from ours in that it imposes specific assumptions about how the agent resolves uncertainty for any given complex contract, namely, the agent is a maxmin decision maker. In contrast, we take a robust approach, and do not impose any assumptions on how agents resolve their confusion in complex mechanisms. Moreover, the uncertainty about the game form in [Jakobsen \(2020\)](#) stems from imperfect memory and limited deductive (computational) ability, while we study uncertainty resulting from strategic interactions.

Our paper is related to the growing literature that studies bounded rationality in industrial organization and shows that firms might benefit from purposefully confusing consumers in the presence of naive consumers; see [Spiegler \(2011\)](#) and references therein. [Glazer and Rubinstein \(2012\)](#) and [Glazer and Rubinstein \(2014\)](#) study persuasion/mechanism design models with boundedly rational agents and show that the listener could benefit from using complex mechanisms that make it difficult for dishonest agents to cheat. These papers impose specific assumptions about the reasoning procedures of the speaker.

2 Preliminaries

Environment. There is a finite set of N agents, $\mathcal{N} = \{1, 2, \dots, N\}$, and an arbitrary (possibly infinite) set of alternatives \mathcal{X} . Each agent i has payoff-relevant information indexed by $\theta_i \in \Theta_i$, where Θ_i is finite. We refer to θ_i as agent i 's type. The set of possible type profiles is $\Theta = \times_{i \in \mathcal{N}} \Theta_i$ with representative element $\theta = (\theta_1, \theta_2, \dots, \theta_N)$. The type profile is distributed according to a prior probability distribution $\pi \in \Delta\Theta$. As is standard, we write $\theta_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$ for a type profile of agents other than agent i .

Each agent i is endowed with a utility function $u_i : \mathcal{X} \times \Theta_i \rightarrow \mathbb{R}$ (we assume private values). The designer has a utility function $v : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$. That is, $u_i(x, \theta_i)$ and $v(x, \theta)$ denote type θ_i 's utility and the designer's utility, respectively, when the type profile is θ and the implemented alternative is x . We assume that the designer is an expected utility maximizer with respect to the distribution π of types. We make no such assumption about the agents because how they form beliefs about other agents' types is irrelevant given the solution concepts we consider.

The designer may wish to use randomization in the mechanism. To incorporate this possibility, we introduce a “dummy” agent $i = 0$ with no preferences, called Nature. The game form is allowed to feature nodes at which Nature is called to play, and the distribution over strategies is picked by the designer as part of the mechanism. We use “bar” to denote the extended profiles that include Nature, for example, $\bar{\mathcal{N}} = \mathcal{N} \cup \{0\}$.

Mechanisms. We consider finite mechanisms that are imperfect-information, extensive-form games with perfect recall and consequences in \mathcal{X} .¹¹ The definition is standard; to shorten the exposition, we only introduce notation associated with a generic game Γ that we are going to use:

- (1) \mathcal{H} is the set of histories, with representative element h , and h_\emptyset denoting the initial (empty) history;
- (2) \subset is the precedence relation over histories;
- (3) \mathcal{Z} is the set of terminal histories, with representative element z ;
- (4) $g(z) \in \mathcal{X}$ is the outcome resulting from z ;
- (5) I_i denotes an information set of agent i ;
- (6) $A(I_i)$ is the set of actions available at information set I_i ;
- (7) A (pure) strategy S_i chooses an action $a \in A(I_i)$ at every information set I_i of agent $i \in \bar{\mathcal{N}}$, and \mathcal{S}_i is the collection of all (pure) strategies for player i ;¹²
- (8) A strategy profile $\bar{S} = (S_0, S_1, \dots, S_N)$ specifies a strategy for each player;
- (9) $z(h, \bar{S})$ denotes the terminal history that results when we start at history h and play proceeds according to the strategy profile \bar{S} ;
- (10) $\pi_0 \in \Delta(\mathcal{S}_0)$ is a full-support probability distribution over Nature’s strategies.

We say that the information set I_i is on the path of play of strategy S_i if there exists \bar{S}_{-i} and $h \in I_i$ such that $h \subset z(h_\emptyset, S_i, \bar{S}_{-i})$. Given two strategies S_i and S'_i , we define $\beta(S_i, S'_i)$ to be the set of information sets that are on the path of play of both S_i and S'_i . Under perfect recall, $I_i \in \beta(S_i, S'_i)$ implies that S_i and S'_i choose the same actions at all

¹¹Finiteness is an important assumption: [Jackson \(1992\)](#) shows that infinite mechanisms may implement virtually any decision rule in undominated strategies, relying on an infinite hierarchy of weakly dominated strategies with no dominant strategy “at the top.”

¹²We focus on pure strategies. This does not affect our negative results that simple mechanisms are (weakly or strongly) dominated. Since a strategy that is not weakly dominated by a pure strategy could nevertheless be dominated by a mixed strategy, allowing mixed strategies gives a weakly smaller set of undominated strategies and thus makes it easier to show that simple mechanisms are dominated.

information sets preceding I_i . If $I_i \in \beta(S_i, S'_i)$ but S_i and S'_i choose different actions at I_i , then we call I_i an earliest point of departure for these two strategies. Following Li (2017), we let $\alpha(S_i, S'_i)$ denote the set of all earliest points of departure for these two strategies. Finally, we let $\Phi(S_i, I_i) = \{S'_i : I_i \in \beta(S_i, S'_i), S_i(I_i) = S'_i(I_i)\}$ be the set of player i 's strategies that agree with S_i at I_i and all information sets that precede it.

Solution concepts. We define three solution concepts that we use in the paper.

Definition 1 (SP). S_i is dominated (or SP-dominated) for type θ_i of agent i if there exists another strategy S'_i such that for all \bar{S}_{-i} ,

$$u_i(g(z(h_\emptyset, S_i, \bar{S}_{-i})), \theta_i) \leq u_i(g(z(h_\emptyset, S'_i, \bar{S}_{-i})), \theta_i),$$

with the inequality being strict for some \bar{S}_{-i} .

Definition 2 (OSP). S_i is obviously dominated (or OSP-dominated) for type θ_i of agent i if there exists another strategy S'_i such that for all $I_i \in \alpha(S_i, S'_i)$,

$$\max_{h \in I_i, \bar{S}_{-i}} u_i(g(z(h, S_i, \bar{S}_{-i})), \theta_i) \leq \min_{h \in I_i, \bar{S}_{-i}} u_i(g(z(h, S'_i, \bar{S}_{-i})), \theta_i),$$

with the inequality being strict for some $I_i \in \alpha(S_i, S'_i)$.

Definition 3 (SOSP). S_i is strongly obviously dominated (or SOSP-dominated) for type θ_i of player i if there exists another strategy S'_i such that for all $I_i \in \alpha(S_i, S'_i)$,

$$\max_{h \in I_i, \bar{S}_{-i}, R_i \in \Phi(S_i, I_i)} u_i(g(z(h, R_i, \bar{S}_{-i})), \theta_i) \leq \min_{h \in I_i, \bar{S}_{-i}, R'_i \in \Phi(S'_i, I_i)} u_i(g(z(h, R'_i, \bar{S}_{-i})), \theta_i),$$

with the inequality being strict for some $I_i \in \alpha(S_i, S'_i)$.

In words, a strategy S_i is dominated for type θ_i if there exists another strategy S'_i that yields a higher payoff for type θ_i for any fixed strategy profile for other players and Nature. A strategy S_i is obviously dominated for type θ_i if there exists another strategy S'_i such that, starting at any earliest point of departure, the worst possible payoff under S'_i for type θ_i across all strategies of other players and Nature is higher than the best possible payoff under S_i for type θ_i across all strategies of other players and Nature. Finally, a strategy S_i is strongly obviously dominated for type θ_i if there exists another strategy S'_i such that, starting at any earliest point of departure, the

worst possible payoff under S'_i for type θ_i across all strategies of other players, Nature, and moves of player i at future information sets is higher than the best possible payoff under S_i for type θ_i across all strategies of other players, Nature, and moves of player i at future information sets. The three dominance concepts are nested: Strongly obviously dominated strategies are obviously dominated, and obviously dominated strategies are dominated. From now on, we fix a dominance concept $K \in \{\text{SP}, \text{OSP}, \text{SOSP}\}$; unless explicitly stated otherwise, all statements should be understood as holding for any such K .

Remark 1 (On randomization). Under the solution concepts analyzed in this paper, one has to take a stance on how agents reason about randomization: When evaluating a strategy, they could either (i) take expectations with respect to designer’s randomization within the mechanism, or (ii) condition on each outcome of the randomization device. Our definitions capture the second possibility. This implies that Definition 1 of dominance for strategy-proof mechanisms is somewhat non-standard.¹³ However, this convention keeps the three solution concepts more analogous, since in OSP (see the baseline model of Li (2017)) and SOSP (see Pycia and Troyan (2021)), dominance is defined with respect to Nature’s moves. Such approach seems natural given the focus on simplicity: We effectively assume that—in the mind of a player—moves of Nature and moves of a strategic player with a constant utility function are treated symmetrically. Since we imposed no a priori restrictions on the space of primitive outcomes \mathcal{X} , it is possible that \mathcal{X} already contains lotteries over deterministic decisions (and the utilities $u_i(x, \theta_i)$ are expected utilities derived from some primitive utility for deterministic decisions); in that case, randomization via Nature’s moves is redundant and our definition of strategy-proof mechanisms reduces to a standard one. Thus, our stronger notion of strategy-proofness has bite only when we explicitly assume that \mathcal{X} does not contain lotteries (e.g., in Section 3.2).

Fixing a game Γ and type θ_i of agent $i \in \mathcal{N}$, we can think of the dominance relation as defining a partial order $\prec_{\theta_i}^K$ on the set of player i ’s strategies: $S_i \prec_{\theta_i}^K S'_i$ if S_i is K -dominated by S'_i for θ_i . We call two strategies S_i and S'_i payoff-equivalent for

¹³We are not first to consider this extension of the standard notion. A (randomized) mechanism is called a universally-truthful mechanism—a common solution concept in the computer science literature, see, for example, Nisan and Ronen (2001)—if every mechanism in the support of the randomized mechanism is a dominant-strategy mechanism.

type θ_i if $u_i(g(z(h_\emptyset, S_i, \bar{S}_{-i})), \theta_i) = u_i(g(z(h_\emptyset, S'_i, \bar{S}_{-i})), \theta_i)$ for all \bar{S}_{-i} . A strategy S_i is K -dominant if all non-payoff-equivalent strategies S'_i are K -dominated by it.

Strategic confusion, simple and complex mechanisms. Throughout, we adopt the following assumption about the rationality of players.

Assumption 1 (Strategic sophistication). *For any $i \in \mathcal{N}$, no type θ_i of player i plays a K -dominated strategy.*

Assumption 1 is the same assumption that forms the basis for the various notions of simplicity that we cover in this paper. For example, if one thinks that participants will play obviously dominant strategies when they are available, then one seems also compelled to reason that participants will avoid obviously dominated strategies (even when no strategy is obviously dominant). From a practical perspective, we can think of the mechanism designer as both designing the game and providing guidance to participants on how they should behave. Assumption 1 sets a lower bound on what recommendations from the designer the players will find persuasive; namely, any player can be persuaded not to play any K -dominated strategy. This is a relatively weak assumption. For example, we do not assume that the players will follow a recommendation whose validity relies on iterative elimination of dominated strategies (from an epistemic perspective, this would require some degree of knowledge of rationality among the players, which we do not assume).

For mechanism Γ with the corresponding set of possible strategies \mathcal{S}_i of player i , we define the set of strategies that type θ_i might play under Assumption 1:

$$U_i^K(\theta_i) = \{S_i \in \mathcal{S}_i : \nexists S'_i \in \mathcal{S}_i, S_i \prec_{\theta_i}^K S'_i\}.$$

That is, $U_i^K(\theta_i)$ is the set of K -undominated strategies for type θ_i . Given the finiteness of the mechanism, that set is non-empty.

If type θ_i has a K -dominant strategy, then the set $U_i^K(\theta_i)$ is not necessarily a singleton; however, all strategies in $U_i^K(\theta_i)$ must be payoff-equivalent for θ_i . In this case, the player is truly indifferent between these dominant strategies, and hence we do not want to label this as “strategic confusion” (otherwise, no solution concept could avoid strategic confusion). It is customary in mechanism design to let the designer select which dominant strategy the player should play, and we follow the same convention

here.¹⁴ Formally, we will treat $U_i^K(\theta_i)$ as consisting of equivalence classes of strategies that are payoff-equivalent for θ_i , and we let the designer pick the representative of each equivalence class (we leave this description verbal in order not to further complicate our notation).

Definition 4 (Strategic confusion and complex mechanism). *Fixing a mechanism Γ , type θ_i of player $i \in \mathcal{N}$ is said to be strategically confused (under solution concept K) if $U_i^K(\theta_i)$ contains at least two (not payoff-equivalent) strategies. In such case, we call mechanism Γ complex (for type θ_i of agent i).*

Strategic confusion means that at least one player has more than one K -undominated strategy in the mechanism, and therefore, in the absence of further assumptions on behavior or strategic reasoning, it is impossible to determine which strategy she will select. Of course, a valid interpretation is that player i “knows” which strategy to choose but the designer is not willing to make an assumption about that choice.

The literature has identified classes of “simple” mechanisms that can be played even if agents are strategically unsophisticated, in the sense of only satisfying Assumption 1. When players cannot forecast their opponents’ play, they will nevertheless know how to behave if they have a dominant strategy (a defining property of SP mechanisms). If, moreover, players cannot engage in contingent reasoning, they should be offered an obviously dominant strategy (leading to the notion of OSP mechanisms, see Li (2017)). Finally, if players lack foresight and cannot predict their own future moves, they should have a strongly obviously dominant strategy (leading to the notion of SOSP mechanisms, see Pycia and Troyan (2021)). This motivates our definition of a simple mechanism.

Definition 5 (Simple mechanism). *A mechanism Γ is simple if for any agent $i \in \mathcal{N}$, no type θ_i is strategically confused.*

Strong and weak dominance of mechanisms. An advantage of a simple mechanism from the point of view of the designer is that she can predict how agents will behave. In contrast, if any agent is strategically confused, the designer—based only on Assumption

¹⁴One justification for this assumption is that the mechanism is paired with “recommended” strategies for the players, and the players follow the recommendation as long as it constitutes a dominant strategy. In many settings the mechanism can be perturbed into a mechanism that is arbitrarily close in terms of payoffs and such that the player is no longer indifferent between these dominant strategies. Finally, the optimal mechanism would often fail to exist in the absence of this assumption.

1—cannot determine the path of play. This seems to provide a strong argument in favor of simple mechanisms. However, that benefit is diminished if the designer can achieve better outcomes using a mechanism that confuses some types.

To formalize this idea, let \mathbf{S}_i denote a type-strategy for player i , that is, $\mathbf{S}_i(\theta_i)$ is the strategy selected by type θ_i of player $i \in \mathcal{N}$. We let $\mathbf{S}_i \subset U_i^K$ mean that \mathbf{S}_i is a selection from the correspondence of K -undominated strategies, i.e., $\mathbf{S}_i(\theta_i) \in U_i^K(\theta_i)$ for all $\theta_i \in \Theta_i$. We define a correspondence

$$V(\Gamma) = \text{Conv. Hull} \left(\left\{ \mathbb{E}_{\theta \sim \pi, S_0 \sim \pi_0} [v(g(z(h_\emptyset, (S_0, \mathbf{S}_1(\theta_1), \dots, \mathbf{S}_N(\theta_N))))), \theta] : \{\mathbf{S}_i \subset U_i^K\}_{i \in \mathcal{N}} \right\} \right)$$

to be the range of the designer's expected payoffs over all possible ways in which confused types can resolve their strategic confusion.¹⁵ By definition, $V(\Gamma)$ is a singleton when Γ is a simple mechanism but it may be a (closed) interval when Γ is complex.

Definition 6 (Strong dominance). *A mechanism Γ is strongly dominated if there exists a mechanism Γ' such that*

$$\max V(\Gamma) < \min V(\Gamma').$$

Definition 7 (Weak dominance). *A mechanism Γ is weakly dominated if there exists a mechanism Γ' such that*

$$\max V(\Gamma) \leq \min V(\Gamma') \text{ and } \max V(\Gamma) < \max V(\Gamma').$$

In words, a mechanism Γ is strongly dominated by a mechanism Γ' if the expected payoff for the designer under Γ' is strictly larger than the expected payoff under Γ , regardless of how confused agents select their strategy. A mechanism Γ is weakly dominated by a mechanism Γ' if the expected payoff for the designer under Γ' is at least as large as the expected payoff under Γ , regardless of how confused agents select their strategies; moreover, the expected payoff under Γ' is strictly larger under some selection. Throughout, we will primarily apply these definitions to the case in which Γ is an *optimal* simple mechanism (that is, it maximizes the designer's expected payoff among all simple mechanisms), in which case it can only be weakly or strongly dominated by a complex mechanism.

¹⁵Note that when computing the range of designer's payoffs, we assume that which undominated strategy a player selects cannot depend on the realization of other players' types.

Participation constraints. Even when simple mechanisms are dominated, a possible argument in their favor is that agents may be discouraged from participation if they face a complex mechanism. However, we also analyze the case in which participation decisions are endogenous. In many settings (including the ones in which the designer’s objective is to maximize her revenue), it is necessary to impose a participation constraint for the problem to be well defined. For a simple mechanism, under the solution concepts in question, it is natural to impose ex-post individual-rationality which requires that for any player i and type $\theta_i \in \Theta_i$, the chosen strategy $S_i = \mathbf{S}_i(\theta_i)$ yields a non-negative payoff in all possible cases:

$$u_i(g(z(h_\emptyset, S_i, \bar{S}_{-i})), \theta_i) \geq 0, \quad \forall \bar{S}_{-i}.^{16} \quad (1)$$

We introduce two extensions of this condition to complex mechanisms. We say that a mechanism Γ provides *partial incentives to participate* if for all $i \in \mathcal{N}$, all $\theta_i \in \Theta_i$, there exists $S_i \in U_i^K(\theta_i)$ such that condition (1) holds. We say that a mechanism Γ provides *full incentives to participate* if for all $i \in \mathcal{N}$, all $\theta_i \in \Theta_i$, for all $S_i \in U_i^K(\theta_i)$ condition (1) holds. The notion of partial incentives to participate is more appropriate when non-participation is thought of as a strategy: The condition states that the non-participation strategy is (obviously) dominated, and hence is not chosen by a rational player. The notion of full incentives to participate is more appropriate if non-participation is thought of as an option for each player to walk away from the mechanism at any point, including after learning her final payoff.

3 Strong dominance

In this section, we study strong dominance of simple mechanisms. Section 3.1 discusses a sufficient condition for the best simple mechanism to be strongly dominated. Section 3.2 shows that strong dominance is possible even in single-agent mechanisms. Section 3.3 provides a condition under which the optimal simple mechanism is not strongly dominated.

¹⁶This implies that strategy S_i obviously dominates non-participation, but not that it strongly obviously dominates non-participation. For SOSF, one may wish to require this condition also for all strategies $R_i \in \Phi(S_i, I_i)$, where I_i is the first information set of player i , for each possible history (this would not change our results in any substantial way).

3.1 Simple mechanisms can be overly restrictive

A weakness of simple mechanisms is that they can be overly restrictive by requiring that *no* type should be confused. In some settings, the presence of certain preference types implies that the set of outcome functions implementable by simple mechanisms is small. However, the behavior of some agents' types could be insignificant for the designer, either because (i) these types do not contribute to the designer's payoff, or (ii) they have low probability. In such cases, the designer may want to impose simplicity only for *a subset of types* with the remaining types being potentially confused.

To formalize this idea, we introduce a property that we call the “accommodation of additional types” (AAT). We show that if the AAT property is violated, then there exists either a payoff function for the designer or a distribution of types such that the best simple mechanism is strongly dominated.

Fixing the type space Θ and players' preferences (the “implementation environment”), we say that an outcome function $\lambda : \Theta \rightarrow \Delta(\mathcal{X})$ is *simply-implementable* if there exists a simple mechanism Γ whose equilibrium leads to the outcome $\lambda(\theta)$ whenever the realized type profile is θ . We let $\Lambda(\Theta)$ denote the set of simply-implementable outcome functions on the type space Θ . We denote by $\Theta \setminus \{\theta_i\}$ the type space $\Theta_1 \times \dots \times (\Theta_i \setminus \{\theta_i\}) \times \dots \times \Theta_N$.

Definition 8. *An implementation environment has the accommodation of additional types (AAT) property if for any i , any $\theta_i \in \Theta_i$, and any outcome function $\lambda \in \Lambda(\Theta \setminus \{\theta_i\})$, there exists $\bar{\lambda} \in \Lambda(\Theta)$ such that $\bar{\lambda}|_{\Theta \setminus \{\theta_i\}} = \lambda$.¹⁷*

The AAT property says that given an arbitrary simple mechanism on the type space $\Theta \setminus \{\theta_i\}$, we can always “accommodate” an additional type θ_i of agent i , that is, assign a K -dominant strategy to θ_i while keeping the outcome of the mechanism for the remaining types unchanged.

By pairing an implementation environment with the designer's objective v (the “design environment”), we can formulate a weaker version of AAT that only applies to designer's payoffs.

Definition 9. *A design environment has the accommodation of additional types (AAT) property if for any i , any $\theta_i \in \Theta_i$, the maximized value of v over $\Lambda(\Theta \setminus \{\theta_i\})$ and over $\Lambda(\Theta)$ is the same under any distribution π supported on $\Theta \setminus \{\theta_i\}$.*

¹⁷We write $\bar{\lambda}|_{\Theta \setminus \{\theta_i\}}$ for the outcome function $\bar{\lambda}$ restricted to $\Theta \setminus \{\theta_i\}$.

Clearly, when AAT fails in some design environment, then it must also fail in the underlying implementation environment. The opposite is not true: failure of AAT in an implementation environment only implies that certain outcome functions are no longer simply-implementable when the incentive constraints of an additional type θ_i must be satisfied, but if maximizing v does not require accessing these outcomes, then AAT will still hold for the design environment.

Proposition 1. *Suppose that the AAT property fails for some (i) implementation environment or (ii) design environment. Then, respectively,*

- (i) for any full-support distribution of types, there exists a payoff function for the designer such that the best simple mechanism is strongly dominated;*
- (ii) there exists a full-support distribution of types such that the best simple mechanism is strongly dominated.*

The intuition behind Proposition 1 is straightforward. Suppose that the AAT property fails in an implementation environment when some type θ_i is “added” to the type space $\Theta_i \setminus \{\theta_i\}$. Then, there exist objective functions for the designer that are maximized at some outcome function that is simply-implementable on $\Theta \setminus \{\theta_i\}$ but not on Θ , and that do not depend “too much” on the outcome implemented for type θ_i . Under such objective functions, the designer can do strictly better by effectively ignoring type θ_i : She offers the same mechanism Γ that she would have offered if type θ_i were not present. The mechanism is complex on Θ because θ_i is strategically confused. However, when the designer’s payoff is not very sensitive to the outcome that occurs conditional on θ_i , Γ will strongly dominate the optimal simple mechanism on Θ . When AAT fails in a design environment, the insignificance of type θ_i for the designer’s overall payoff can be achieved by choosing a distribution of types π under which θ_i is very improbable. The proof of Proposition 1—which formalizes these arguments—can be found in Appendix A.2.

Example 3. Consider a voting environment with two agents and three alternatives, $\mathcal{X} = \{a, b, c\}$. Each agent’s type can be represented as a ranking of the three alternatives. More specifically, each agent gets utility 1 if her top choice is implemented, 1/2 if her second choice is implemented, and 0 otherwise. The distribution of types π is i.i.d. uniform. The designer would like to maximize welfare but is Rawlsian and risk-averse: If u_i is the

ex-post utility of agent i , then the designer's payoff is $v(\min(u_1, u_2))$ for some strictly concave and increasing function v .

The best simple mechanism (for any $K \in \{\text{SP}, \text{OSP}, \text{SOSP}\}$) is dictatorship with full range.¹⁸ The outcome function of that mechanism (with the row player being a dictator) is illustrated in Table 1.

Table 1: The best simple mechanism

	<i>abc</i>	<i>acb</i>	<i>bac</i>	<i>bca</i>	<i>cab</i>	<i>cba</i>
<i>abc</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>acb</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>bac</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
<i>bca</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
<i>cab</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>
<i>cba</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>	<i>c</i>

Table 2: A complex mechanism

	<i>abc</i>	<i>acb</i>	<i>bac</i>	<i>bca</i>	<i>cab</i>	<i>cba</i>
<i>abc</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>acb</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>bac</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
<i>bca</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>
<i>cab</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>c</i>	<i>c</i>	<i>c</i>
<i>cba(1)</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>c</i>	<i>c</i>	<i>c</i>
<i>cba(2)</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>	<i>b</i>

Consider instead a delegation mechanism in which the designer specifies a set of menus from which the dictator chooses one, with the second player selecting an alternative from the chosen menu. The dictator can choose menu $\{a\}$, menu $\{b\}$, or menu $\{a, c\}$.

Types abc, acb, bac, bca of the row player retain their K -dominant strategies. Type cab has a K -dominant strategy to choose the menu $\{a, c\}$. However, type cba does not have a K -dominant strategy: Both menu $\{b\}$ and menu $\{a, c\}$ are not K -dominated—the mechanism is complex. This is illustrated in Table 2.

However, whichever strategy type cba chooses, the expected payoff to the designer is strictly higher than in the best simple mechanism. Indeed, consider first the case in which type cba chooses menu $\{a, c\}$. Then, the difference in expected payoffs to the designer between the complex mechanism and the best simple mechanism (which can be calculated by comparing the cells of the two tables with different outcomes) is $\frac{1}{36} \times \left[v\left(\frac{1}{2}\right) - v(0) \right] > 0$. Now, consider the case in which type cba chooses menu b . The difference in the expected payoffs of the designer is $\frac{1}{36} \times \left[-2v(1) + 4v(1/2) - 2v(0) \right] > 0$, by strict concavity of v . Thus, the complex mechanism is guaranteed to yield a strictly higher expected payoff to the designer regardless of how type cba resolves her confusion.

Note that the AAT property is violated in the above implementation environment:

¹⁸It suffices to show this for the concept of SP since it is the most permissive one (and because dictatorship is a simple mechanism under all notions). Due to our Definition 1 of SP, it is without loss of optimality to look at deterministic mechanisms, and thus only ordinal preferences of players matter. We can directly verify that dictatorship with full range is optimal for our objective function.

The simple mechanism defined by the first 5 rows of Table 2 by excluding type cba cannot be extended to a simple mechanism with type cba added back. To see that, note that for type cba to have a K -dominant strategy, she must be able to secure c conditional on $\theta_2 = bca$; she must also be able to secure b or c conditional on $\theta_2 = bac$. If type cba gets c in any contingency conditional on $\theta_2 = bac$, then type cab of player 1 no longer has a K -dominant strategy. Thus, type cba must get b conditional on $\theta_2 = bac$. But in that case, type bca of player 2 no longer has a K -dominant strategy. While type cba influences the designer’s payoff, the designer nevertheless prefers to “ignore” that type by implementing a mechanism that is simple for the remaining types and that makes cba strategically confused. \square

More generally, the AAT property fails in many classical social choice environments. It is well known that the only SP mechanisms whose range contains at least three alternatives are dictatorships; however, there are nontrivial strategy-proof social choice functions on restricted domains. Indeed, much of the research on strategy-proof social choice can be seen as establishing possibility results for (various) restricted domains.¹⁹ Thus, the AAT property fails. Proposition 1 then implies that in the social choice environment, it might be beneficial to “ignore” some types and employ a social choice rule that is SP-implementable on a smaller domain, rather than using a SP mechanism on a larger domain. In particular, this will hold if the “problematic” types occur with sufficiently low probability. For example, it is known that a social choice function on profiles of single-peaked preferences over a totally ordered set is strategy-proof if and only if it is a generalized median voter scheme. If the designer finds the outcome function of some generalized median voter scheme more desirable than that of a dictatorship, and the probability that the true type profile is contained in the set of single-peaked preferences is high enough, then the generalized median voter scheme strongly dominates the best simple mechanism on the full domain. Similar examples can be found for OSP and SOSp.²⁰

Note that AAT holds in any environment with a single agent. The result is immediate: A single agent in a simple mechanism effectively chooses her most preferred outcome from some (potentially stochastic) menu.

¹⁹We refer interested readers to Barberà (2010) for a survey on strategy-proof social choice.

²⁰See Arrillaga et al. (2020) and Bade and Gonczarowski (2017) who characterize the class of OSP-implementable and unanimous social choice functions for single-peaked preferences, among other applications.

Claim 1. *The AAT property holds in any environment with $N = 1$.*

3.2 Strong dominance with a single agent

Claim 1 in the previous subsection showed that the best simple mechanism with a single agent cannot be strongly dominated due to failure of the AAT property. Moreover, in the absence of any opponents, it may seem that the agent’s strategic problem is trivial. Nevertheless, we prove that strong dominance is possible even in that case.²¹

Claim 2. *The best simple mechanism may be strongly dominated in the case $N = 1$.*

Clearly, the conclusion of Claim 2 is only possible when the agent is not assumed to be Bayesian. Instead, as captured by our definition of SP, OSP, and SOSp, the agent compares different strategies by conditioning on every realization of the randomization device used by the designer (see also Remark 1). Here, we prove the claim for the solution concept of SP, since the example of strong dominance is particularly easy to construct; in Appendix A.4, we construct a more complicated example that delivers the conclusion for all three solution concepts.

Example 4 (Proof of Claim 2 for SP). The single agent has two possible types, u and d . Let π_u and $\pi_d = 1 - \pi_u$ denote the respective probabilities of these two types. There are 4 alternatives; $\mathcal{X} = \{U, U', D, D'\}$. The preferences of the types are given by any cardinal utility with the following ordinal consequences:

1. type u : $U' > D > U > D'$;
2. type d : $D' > U' > D > U$.

The designer receives a payoff of 1 if the type is $j \in \{u, d\}$ and the outcome is J ; she receives 1/2 when the type is j and the outcome is J' ; in all other cases, she receives 0.

Consider the following mechanism Γ in which the designer uses a randomization device that has two equally likely outcomes “Heads” and “Tails”, with four possible strategies $(s_u^H, s_u^T, s_d^H, s_d^T)$:

²¹Tillio et al. (2017) study the sale of an object to an ambiguity averse buyer, and show that the seller can increase her profit by using an ambiguous mechanism. Their analysis—unlike ours—hinges on the assumption that the buyer is ambiguity averse.

	“Heads”	“Tails”
s_u^H	U	U'
s_u^T	U'	U
s_d^H	D	D'
s_d^T	D'	D

This mechanism has a simple interpretation: The player is asked to choose between a 50 – 50 lottery over U and U' , and a 50 – 50 lottery over D and D' . However, the player must additionally choose the side of the coin leading to her preferred outcome.

The mechanism Γ is not SP according to Definition 1: type u is confused between strategies s_u^H and s_u^T , and type d is confused between strategies s_d^H and s_d^T . Put differently, the agent is able to recognize the inferior lottery as dominated, but she is confused about which side of the coin to “bet on.” From the perspective of the designer, however, that confusion is irrelevant, since the choice of the side of the coin does not affect the marginal probabilities of outcomes: The mechanism generates an expected payoff of $3/4$ for the designer regardless of how the agent resolves her confusion, and regardless of the prior distribution of types.

It is not possible to implement the same outcome (a 50 – 50 lottery over U and U' conditional on u , and a 50 – 50 lottery over D and D' conditional on d) with a simple mechanism. We prove this by showing that when $\pi_d = 2/3$, the *best* SP mechanism is to always implement D which yields an expected payoff of only $2/3$. This will also prove Claim 2 (for the case of SP).

We only have to prove that the designer cannot do better than $2/3$ with a simple mechanism when $\pi_d = 2/3$. Note that randomization cannot increase the designer’s payoff in the best simple mechanism. Indeed, for each SP mechanism in which the designer uses a randomization device, the designer can do weakly better by always selecting the same outcome of the randomization device; namely the one associated with the highest conditional expected payoff. Thus, it is without loss of generality to restrict attention to deterministic SP mechanisms. And since there is a single agent, any deterministic SP mechanism is equivalent to a menu of deterministic outcomes for the agent to choose from. To get an expected payoff of more than $2/3$, the designer must offer D in the mechanism, and D must be chosen by type d (indeed, in the opposite case, the designer’s expected payoff is upper-bounded by $1/3 + 2/3 \cdot 1/2 = 2/3$). However, in this case, neither D' nor U' can be offered, as then d would not choose D . Hence, either

only D is offered, in which case the designer gets an expected payoff of $2/3$, or D and U are offered, in which case type u also selects D , and the expected payoff to the designer is again $2/3$. \square

Summarizing the example, the *only* way to implement a 50 – 50 lottery over U and U' conditional on u , and a 50 – 50 lottery over D and D' conditional on d is by creating strategic confusion (through the agent’s choice of the side of the coin to bet on). To understand the construction of the complex mechanism, note that the dominance relationship forms a cycle: s_u^T dominates s_d^H for type u , s_d^H dominates s_u^H for type d , s_u^H dominates s_d^T for type u , and s_d^T dominates s_u^T for type d . However, no other pair of strategies are comparable, and thus no undominated strategy for any type can be dispensed with. For example, suppose that the designer uses a mechanism that only offers strategies s_u^H and s_d^H , that is, she predetermines the side of the coin leading to each given outcome. Then, type u will be confused, and hence she might play s_d^H , leading to the lottery over D and D' being implemented for both types.

We now prove a general result that exposes the underlying logic of the complex mechanism in Example 4, and provides a tractable way of constructing a superior complex mechanism. We first define the notion of ordinal incentive compatibility. Let $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$, and suppose that there are m alternatives, $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$. An ordinal incentive-compatible allocation is a mapping from preference orderings to stochastic outcomes $f : \Theta \rightarrow \Delta\mathcal{X}$ such that for any $\theta, \theta' \in \Theta$, $f(\theta)$ first-order stochastically dominates $f(\theta')$ with respect to the preference ordering θ , that is, for all $x \in \mathcal{X}$,

$$\sum_{y \in \mathcal{X}: u(y, \theta) \geq u(x, \theta)} f(\theta)(y) \geq \sum_{y \in \mathcal{X}: u(y, \theta) \geq u(x, \theta)} f(\theta')(y),$$

where $f(\theta)(y)$ denotes the probability of outcome y under distribution $f(\theta)$.²²

Proposition 2. *Let $N = 1$ and $K = SP$. If there exists an ordinal incentive-compatible allocation that yields a strictly higher expected payoff for the designer than the best deterministic mechanism, then the best simple mechanism is strongly dominated.*²³

The proof of Proposition 2 is constructive (see Appendix A.5). We show that for any ordinal incentive-compatible allocation, there exists a perturbation of that allocation

²²Even though the definition refers to the cardinal representation assumed in Section 2, note that it only uses ordinal information about preferences.

²³It is natural to conjecture that the converse might also be true, but it is not. As this is orthogonal to the message of the paper, we do not include the counterexample in this draft.

that can be implemented by a complex mechanism. If the optimal ordinal incentive-compatible allocation is strictly better for the designer than the best deterministic mechanism, then the conclusion follows from the observation that there always exists an optimal simple mechanism that is deterministic. Note that the random allocation implemented by the complex mechanism in Example 4 indeed satisfies ordinal incentive compatibility.

3.3 An optimality foundation for simple mechanisms

In this subsection, we show that in a class of environments that we define below, the best simple mechanism is not strongly dominated. The key observation (extending an insight from Yamashita (2015a)) is that for any finite mechanism Γ (simple or complex), $\min V(\Gamma)$ is weakly less than the maximum expected payoff the designer could obtain in a mechanism that only satisfies a subset of incentive constraints that correspond to the edges of an arbitrary tree in the type space. Thus, if the designer's expected payoff from the best simple mechanism is the same as that in the relaxed problem with incentive constraints that correspond to the edges of some tree in the type space, then the best simple mechanism is not strongly dominated.

Let $G = (V, E)$ be a directed graph with vertex set V and edge set $E \subseteq V \times V$. A graph is called a (rooted) tree if it has exactly one vertex with no outgoing edges (called the "root") and exactly one path from every vertex to the root. For each agent i , consider a tree $T_i = (\Theta_i, E_i)$ where $E_i \subseteq \Theta_i \times \Theta_i$. Each directed edge $(\theta_i, \theta'_i) \in E$, also denoted $\theta_i \rightarrow \theta'_i$, corresponds to the incentive constraint that type θ_i does not want to adopt the strategy of type θ'_i . For each agent i , fix a tree T_i . The collection of trees $\{T_i\}_{i \in \mathcal{N}}$ then defines a relaxed optimization problem in which the only incentive constraints are the ones that correspond to the edges on the trees. Formally, the IC-relaxed problem is to find a mechanism Γ that assigns a strategy $\mathbf{S}_i(\theta_i)$ to each type θ_i of each player $i \in \mathcal{N}$ and maximizes the designer's objective functions among all such mechanisms with the property that if $\theta_i \rightarrow \theta'_i$, then $\mathbf{S}_i(\theta'_i) \prec_{\theta'_i}^K \mathbf{S}_i(\theta_i)$. We ignore participation constraints for now.

Proposition 3. *Suppose that there exists a collection of trees $\{T_i\}_{i \in \mathcal{N}}$ such that the designer's expected payoff from the IC-relaxed optimization problem corresponding to $\{T_i\}_{i \in \mathcal{N}}$ is the same as from using the optimal simple mechanism. Then, the optimal*

simple mechanism is not strongly dominated.

We now explain how to incorporate participation constraints. In the IR-non-relaxed problem, the IR constraint (1) is imposed for every type of every player. In the IR-relaxed problem, the IR constraint (1) is only imposed for the root of each tree T_i .

Proposition 3’. *Suppose that there exists a collection of trees $\{T_i\}_{i \in \mathcal{N}}$ such that the designer’s expected payoff from the (i) IR-relaxed and IC-relaxed or (ii) IR-non-relaxed and IC-relaxed optimization problem corresponding to $\{T_i\}_{i \in \mathcal{N}}$ is the same as from using the optimal simple mechanism with participation constraints. Then, the optimal simple mechanism is not strongly dominated by any mechanism with, respectively, (i) partial or (ii) full incentives to participate.*

Proposition 3(3’) is relatively abstract and does not offer guidance on how to look for the “right” tree in the type space. For illustration, we apply Proposition 3’ to show that the optimal SP mechanism is not strongly dominated in many canonical settings in which the designer maximizes revenue. Since for any mechanism, a strategy that is not dominated is also not (S)OSP-dominated, whenever the optimal SP mechanism can be implemented via an (S)OSP mechanism, our analysis below implies that also the optimal (S)OSP mechanism is not strongly dominated.

We will show that the assumption (i) of Proposition 3’ is satisfied when the uniform shortest-path tree condition holds and the distribution π is regular, as defined by [Chen and Li \(2018\)](#). Loosely speaking, these conditions ensure that, in the optimal SP mechanism, the set of binding constraints for agent i is independent of the types of agents other than i . To be rigorous and self contained, we formally define these terms in the Supplemental Material [OA.1](#). Here, we note that the uniform shortest-path tree condition is of interest because a number of resource allocation problems satisfy it. First, it is satisfied in environments with one-dimensional types and single crossing. This fits many classical applications of mechanism design, including single-unit auctions (e.g., [Myerson \(1981\)](#)), public goods (e.g., [Mailath and Postlewaite \(1990\)](#)), and standard bilateral trade (e.g., [Myerson and Satterthwaite \(1983\)](#)). The uniform shortest-path tree condition also holds in multi-unit auctions with homogeneous or heterogeneous goods, as long as the agents’ private values are one-dimensional. Second, the uniform shortest-path tree condition can also be satisfied in some multi-dimensional environments, such as the auction for capacity constrained bidders (see [Malakhov and Vohra \(2009\)](#)).

By the definition of the uniform shortest-path tree, for each agent, there exists a tree T_i such that the designer’s expected payoff from using the best SP mechanism is the same as the IC-relaxed and IR-relaxed optimization problem corresponding to $\{T_i\}_{i \in \mathcal{N}}$. Thus, Proposition 3’ applies and we obtain the following corollary.

Corollary 1. *In environments in which the uniform shortest-path tree condition holds and π is regular, the optimal SP mechanism is not strongly dominated by any mechanism with partial incentives to participate (hence, it is also not strongly dominated by any mechanism with full incentives to participate).*

Example 1 in the Introduction shows that when the uniform shortest-path tree condition is violated, the conclusion of Corollary 1 can also fail.

Example 5 (Example 1 revisited). Recall that the best simple (OSP) mechanism gives the designer an expected profit of $1/5$. It can be directly verified (by considering all possible trees for both players, see Appendix A.1) that the trees that yield the lowest value of the relaxed problem are $T_A = \{0, 2/3\}$, $0 \rightarrow 2/3$ and $T_B = \{1/3, 1\}$, $1 \rightarrow 1/3$, that is, for trader A we only impose the IC constraint that type 0 does not want to imitate type $2/3$, and for trader B that type 1 does not want to imitate type $1/3$. The value of the IC-relaxed and IR-non-relaxed problem is $1/5$. Thus, the best simple mechanism is *not* strongly dominated by any mechanism with *full* incentives to participate, by Proposition 3’.

However, we already know from Example 1 that the best simple mechanism is strongly dominated by a complex mechanism with *partial* incentives to participate. Indeed, the value of the IC-relaxed and IR-relaxed problem is $2/3$ (under the relaxation that yields the minimal value). This is due to the failure of the uniform shortest-path tree condition: For trader B , even when the low type’s IR constraint holds, and the high type’s IC constraint against deviation to the low type holds, this does not guarantee that the high type’s IR constraint will be satisfied in the optimal mechanism for the relaxed problem.

The superior complex mechanism from Example 1 can be understood as partially capturing the benefits of the relaxed IR constraint of type 1 of trader B . Indeed, if we removed the first offer to trader B from the mechanism, type $1/3$ ’s optimal choice would be unaffected (since accepting the first offer was obviously dominated for that type anyway), while type 1 would no longer be confused—she would follow the same

strategy as type 1/3. However, type 1 would lose its (partial) incentives to participate since following 1/3's strategy sometimes results in a strictly negative payoff. This reveals that the reason for making the first offer to trader B in the complex mechanism is to ensure participation. By confusing type 1 of trader B , the designer extracts more surplus with one undominated strategy (“reject the first offer but accept the next offer”) while dominating non-participation with a second undominated strategy (“accept the first offer”). \square

More generally, failure of the premise of Proposition 3 or 3' indicates that a strict improvement could be made in the problem of finding the best simple mechanism if some IC or IR constraint were relaxed. In a simple mechanism, however, a strategy assigned to each type must satisfy all these constraints, and hence it is not possible to achieve that potential improvement. In a complex mechanism, by contrast, each type may be assigned multiple (undominated) strategies and, at least in principle, no single strategy must satisfy all the constraints. This additional flexibility of complex mechanism may sometimes allow the designer to achieve the improvement.

4 Weak dominance

In this section, we study the concept of *weak dominance*. We demonstrate that weak dominance of simple mechanisms by complex ones is a common occurrence. We first prove a result (Proposition 4) providing conditions under which a given simple mechanism is weakly dominated. We then show how to apply the proposition in the context of revenue maximization; we show that in a fairly large class of problems the best simple mechanism is weakly dominated. We also provide one positive result: Under additional conditions, the posted price mechanism is weakly undominated for the solution concepts of SP and OSP (but not SOS). We provide additional applications of Proposition 4 to efficiency maximization in bilateral trade and a social choice problem without transfers in Supplemental Materials OA.3 and OA.4, respectively.²⁴

²⁴The message of this section is reinforced by additional examples of weak dominance described in the independent work of Mukherjee et al. (2020), as discussed in Section 1.

4.1 Sufficient condition for weak dominance

We begin by formulating an abstract result providing sufficient conditions under which a given simple mechanism Γ is weakly dominated by a complex mechanism. Let $\mathcal{Y} \subseteq \mathcal{X}$ and define

$$\Theta_i^{\mathcal{Y}} = \left\{ \theta_i \in \Theta_i : \max_{x \in \mathcal{Y}} u_i(x, \theta_i) > \min_{\bar{S}_{-i}} u_i(g(z(h_{\emptyset}, \mathbf{S}_i(\theta_i), \bar{S}_{-i})), \theta_i) \right\}.$$

That is, $\Theta_i^{\mathcal{Y}}$ is the set of types of agent i that strictly prefer some outcome in \mathcal{Y} to the worst possible outcome in the simple mechanism Γ .

Proposition 4. *Fix a simple mechanism Γ . Suppose that for any agent $i \in \mathcal{N}$, there exists $\mathcal{Y} \subseteq \mathcal{X}$ and a simple mechanism $\Gamma_{-i}^{\mathcal{Y}}$ played by agents $-i$ with an outcome space \mathcal{Y} such that*

1. *Each outcome $x \in \mathcal{Y}$ occurs at some terminal node in $\Gamma_{-i}^{\mathcal{Y}}$;*
2. *For any $\theta_i \in \Theta_i^{\mathcal{Y}} \neq \emptyset$, the designer prefers (strictly for some $\theta_i \in \Theta_i^{\mathcal{Y}}$) the conditional expected payoff from the mechanism $\Gamma_{-i}^{\mathcal{Y}}$ to the conditional expected payoff from the mechanism Γ .*

Then, the mechanism Γ is weakly dominated.

The proof can be found in Appendix A.7. The idea is straightforward: Given some initial simple mechanism Γ , agent i is offered an additional option that guarantees herself an outcome in \mathcal{Y} . If this option is not chosen, play proceeds as in the original mechanism Γ . The key difference between a player and the designer when evaluating the additional option is that the player is only assumed to avoid K -dominated strategies while the designer is an expected-payoff maximizer. The agent will not rule out a strategy that gives her a high enough payoff in some scenario, no matter how improbable that scenario is. The designer can structure the set \mathcal{Y} and the mechanism $\Gamma_{-i}^{\mathcal{Y}}$ in a way that gives her a high payoff on average, while guaranteeing at least one contingency with a good outcome for player i .

When $K \in \{\text{SP}, \text{OSP}\}$, the above argument implies that for the conclusion of Proposition 4 to hold we only need the assumption to be satisfied for *some* player i (this claim is formalized by the proof in Appendix A.7). However, for the solution concept of SOSPP, the construction might fail because agent i evaluates the payoffs from the original game Γ differently when she has an additional node in the modified game

(corresponding to being offered the additional option). Hence, to cover this case, we need the prerequisite to hold for every agent i ; this allows us to add the additional option to the node of the player who moves first in the original game Γ .

The above construction bears some high-level resemblance to ideas first considered by [Börger and Smith \(2012\)](#) and [Börger \(2017\)](#) in the context of specific design environments. Given an optimal dominant-strategy mechanism, they introduce additional options for players in such a way that the Bayes Nash equilibria of the resulting new mechanism provide weakly higher payoffs to the designer regardless of players' higher-order beliefs. The intuition behind [Proposition 4](#) is indeed similar; however, [Proposition 4](#) applies to different solution concepts and does not rely on fixing any particular environment or objective function.

4.2 Revenue maximization

To show how [Proposition 4](#) can be applied, we analyze a standard revenue-maximization problem with quasi-linear utilities ([Supplemental Materials OA.3](#) and [OA.4](#) contain additional illustrations beyond revenue-maximization and beyond settings with monetary transfers). Let X be the space of possible allocations (which could involve randomization), and define $\mathcal{X} = X \times \mathbb{R}^N$, with $(y, t_1, \dots, t_N) \in \mathcal{X}$ interpreted as an outcome in which allocation y is implemented, and player i pays the designer t_i . We have $u_i((y, t_1, \dots, t_N), \theta_i) = \tilde{u}_i(y, \theta_i) - t_i$, for some arbitrary $\tilde{u}_i(y, \theta_i)$ assumed non-negative and non-constant in θ_i . The designer maximizes expected revenue. We will consider two cases. In the first case, the designer is satisfied with a mechanism that provides partial incentives to participate; we will show that this leads to the best simple mechanism being weakly dominated in a particularly blatant way, with revenue that is unbounded in the best case for the designer. In the second case, the superior mechanism will provide full incentives to participate.

Claim 3. *The revenue-maximizing simple mechanism is weakly dominated by a complex mechanism with partial incentives to participate.*

The superior complex mechanism that we construct in [Appendix A.8](#) exploits—in a very stark way—the possibility that agents lack strategic sophistication. The designer approaches some agent i and proposes to her the following additional bet: Agent i gets a large amount of money M if a coin is heads but pays the designer M if the coin is tails.

The designer biases the coin so that tails has probability $1 - \epsilon$ for an arbitrarily small ϵ . Under our assumed solution concepts, the agent finds accepting the bet K -undominated because she evaluates her outcomes conditional on each realization of the randomization device: When M is sufficiently large, taking the bet dominates all other strategies conditional on the coin being heads. If the agent refuses to take the bet, the designer is still guaranteed the same revenue as in the original simple mechanism; and if the agent accepts the bet, the designer’s expected revenue is unbounded.

It is known (see, for example, [Ashlagi and Gonczarowski \(2018\)](#) and [Pycia and Troyan \(2021\)](#)) that randomization cannot increase the designer’s payoff within the class of simple mechanisms.²⁵ Interestingly, the above example shows that this does not imply that the designer should never randomize when facing unsophisticated agents. On the contrary, randomization can be used to purposefully confuse the agents in order to obtain a superior outcome, at least in some cases.

Even if agents were assumed to be Bayesian with respect to the designer’s randomization, the designer could still achieve an unbounded revenue in the best case of a complex mechanism when $N \geq 2$. The additional option for agent i has the same two outcomes for i , but the choice between them is made by some other agent j who is paid $\epsilon > 0$ to choose the outcome in which agent i pays M . Accepting this modified “bet” is K -undominated for i because agent i is not assumed to believe that agent j will never play a dominated strategy (agents do not engage in iterative elimination of dominated strategies).

The above weakly-dominating mechanisms may seem unlikely to “work” in practice, in the sense that the additional strategic option offered in the complex mechanism is clearly “unattractive” for the agent. We offer two comments: First, similarly to how many results are interpreted in mechanism design, we view the value of these examples as illustrating the possibility that a simple mechanism may be dominated. Our construction is optimized for making the mathematical argument concise; there may exist more subtle ways to weakly dominate a simple mechanism. Second, anecdotally, sellers frequently offer seemingly unattractive options to customers hoping to exploit their potential inability to rank these options as inferior.²⁶ And in any case, the designer never loses by switching to a superior complex mechanism, so she may prefer the complex

²⁵For SP mechanisms, this is a consequence of our Definition 1, see Remark 1.

²⁶See, for example, [Chernev et al. \(2015\)](#) for a survey of results on choice overload.

mechanism even if she thinks that the agent is unlikely to choose the additional option.

A different criticism of the above complex mechanisms is that the agent will sometimes walk away with a negative payoff. We thus turn attention to the case when the complex mechanism is required to provide full incentives to participate. To simplify exposition, we consider the classical problem of allocating a single object to one of N ex-ante identical bidders.²⁷

Claim 4. *Suppose that $N \geq 2$, and let \bar{u} be the highest possible valuation for the object. If it is not a revenue-maximizing simple mechanism to sequentially offer the object at a price of \bar{u} to all the players, then the best simple mechanism is weakly dominated by a complex mechanism with full incentives to participate.*

The proof (see Appendix A.9) is based on Proposition 4 and resembles the mechanism we described previously for the case $N \geq 2$. The main difference is that the additional option is carefully constructed so that it always yields a non-negative payoff to the types that may choose it.

The assumption of Claim 4 cannot be completely relaxed. For a simple example, note that if all players have a value of \bar{u} for the object with probability one, then sequentially offering the object at a price of \bar{u} to all the players is an optimal simple mechanism that is *not* weakly dominated under full incentives to participate (for it to be dominated, there would have to exist an on-path history in which some player is charged more than \bar{u} but that would be incompatible with full incentives to participate). Next, we prove a general positive result for the case of $N = 1$.

Claim 5. *Suppose that the designer sells a single indivisible object to a single agent, attempting to maximize revenue. Assume that there exists a unique optimal simple mechanism (in which case it must be outcome-equivalent to a posted price mechanism). Then, for $K = SP$ and $K = OSP$, that mechanism is not weakly dominated by any complex mechanism with full incentives to participate.*

We emphasize that *not* being weakly dominated is a particularly strong optimality foundation; it implies that any complex mechanism that sometimes leads to a higher payoff for the designer must necessarily be strictly worse in some other case.

²⁷Note that the ascending clock auction with jump-bidding in Example 2 in the introduction provides full incentives to participate under OSP. This is because if type $v \geq b$ makes a jump bid to b at price $p < b$ (jump-bidding is dominated for all other types), then she gets a non-negative payoff after every possible history. Obviously, her payoff is also non-negative if she does not jump-bid.

The proof of Claim 5 is relatively involved (see Appendix A.10). Both assumptions of Claim 5—that the optimal simple mechanism is unique (which holds generically) and that the solution concept is SP or OSP—are needed. In the Supplemental Material OA.2, we construct an example in which the optimal simple mechanism is not unique and we show that it is weakly dominated. Here, we show that the posted price mechanism is weakly dominated under SOSP, except in degenerate cases.

Let p^* be the revenue-maximizing price under π in the optimal SOSP mechanism. Suppose that p^* is not equal to the highest possible type under π . Then, a one-person descending clock auction with a minimum price p^* weakly dominates the posted price mechanism while providing full incentives to participate. This design is formally a sequential posted price mechanism in which the designer starts from a price equal to the highest type and drops the price by some sufficiently small ϵ in every round, with the last price equal to p^* . In each round, the agent accepts or rejects an offer; rejecting the final offer p^* leads to no trade. All types $\theta < p^*$ have an SOSP-dominant strategy to never buy (this strategy gives a non-negative payoff). For types $\theta \geq p^*$, it is SOSP-dominant to reject any price weakly above θ , and SOSP-undominated to buy at any price strictly below θ (again, in all cases, the payoff is non-negative). What makes that last choice undominated for the agent is that when she compares buying at a high price $p \in (p^*, \theta)$ to rejecting p and waiting for a lower offer, she considers the worst-case scenario over her own future moves under SOSP; the worst case is that she will reject all offers and end up without the good, which is strictly worse than buying at p . The designer can only be better off, since trade prices are weakly above p^* ; she is strictly better off whenever the highest type decides to buy at a price strictly between p^* and her value. Thus, the posted price mechanism is weakly dominated.

5 Conclusion

In mechanism design, it seems useful to distinguish (simple) mechanisms in which agents face a straightforward choice problem from (complex) mechanisms that require agents to engage in complex mental tasks to determine their optimal strategy. The literature has made a great deal of progress in terms of formulating different notions of simplicity and characterizing mechanisms that are simple according to these notions. However, the understanding of the design of mechanisms with unsophisticated agents, as we argued

in this paper, is far from complete. Indeed, in many cases, the designer might prefer mechanisms that are not simple, even under the assumption that agents choose strategies that are the worst possible for the designer whenever they are confused.

We suggest some directions for further research. An important avenue is the optimal design of mechanisms when agents are strategically unsophisticated. Our analysis indicates that the optimal design of mechanisms with unsophisticated agents could be challenging: One should not simply optimize over the class of simple mechanisms, at least not without first establishing their optimality. Searching over all mechanisms present new challenges, as the designer cannot rely on any standard revelation principle when mechanisms are evaluated by their worst-case performance.

Relatedly, while we focused primarily on negative results throughout the paper, we expect that establishing optimality foundations for simple mechanisms—not being weakly or strongly dominated—might be a particularly promising research direction. This is because such a foundation can often be found by first solving an easier relaxed problem, and then showing that the upper bound is achieved by a simple mechanism.

While we only considered the strategic dimension and worked with SP, OSP, and SOSP mechanisms in the private-value setting, the same exercise could be done for other notions of simplicity along the strategic dimension (such as the notion of “strategically simple mechanisms” defined by [Börgers and Li \(2019\)](#)), other dimensions of simplicity such as computational complexity, and interdependent-value settings.²⁸

Finally, it would be interesting to conduct experimental tests of the best simple mechanism against the complex mechanism that dominates it, such as the ascending-clock auction and the ascending-clock auction with jump bidding. Real-life settings would allow for factors not accounted for in our theoretical analysis, such as, for example, the possibility that agents would refuse to participate in mechanisms requiring complicated mental calculations. These findings could then be used to support or invalidate the superiority of complex mechanisms.

²⁸In the interdependent-value setting, [Jehiel et al. \(2006\)](#) and [Yamashita and Zhu \(2020\)](#) consider a designer who does not have realizable information about the agents’ beliefs, while the agents are still assumed to play a Bayesian equilibrium. While the focus of these papers differs from ours, we note that Example 5.1 in [Jehiel et al. \(2006\)](#) and the analysis in [Yamashita and Zhu \(2020\)](#) could be used to show that the designer might prefer a mechanism that is not ex post incentive compatible to the optimal ex post incentive-compatible mechanism in the interdependent-value setting, under the assumption that agents do not play weakly dominated strategies. [Yamashita \(2015a\)](#) shows that for revenue maximization in an interdependent-value auction, under certain conditions, a version of a second-price auction (which is neither dominant-strategy nor ex post incentive compatible) is optimal for implementation in weakly undominated strategies.

References

- ARRIBILLAGA, R. P., J. MASSÓ, AND A. NEME (2020): “On Obvious Strategy-Proofness and Single-Peakedness,” *Journal of Economic Theory*, 186, 104992.
- ASHLAGI, I. AND Y. A. GONCZAROWSKI (2018): “Stable Matching Mechanisms are not Obviously Strategy-Proof,” *Journal of Economic Theory*, 177, 405–425.
- BADE, S. AND Y. A. GONCZAROWSKI (2017): “Gibbard-Satterthwaite Success Stories and Obvious Strategyproofness,” in *Proceedings of the 2017 ACM Conference on Economics and Computation*, 565–565.
- BARBERÀ, S. (2010): “Strategy-Proof Social Choice,” in *Handbook of Social Choice and Welfare*, ed. by K. Arrow, A. Sen, and K. Suzumura, North-Holland, vol. 2, 731–831.
- BERGEMANN, D. AND S. MORRIS (2005): “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- BÖRGERS, T. (1991): “Undominated Strategies and Coordination in Normalform Games,” *Social Choice and Welfare*, 8, 65–78.
- (2017): “(No) Foundations of Dominant-Strategy Mechanisms: A Comment on Chung and Ely (2007),” *Review of Economic Design*, 21, 73–82.
- BÖRGERS, T. AND J. LI (2019): “Strategically Simple Mechanisms,” *Econometrica*, 87, 2003–2035.
- BÖRGERS, T. AND D. SMITH (2012): “Robustly Ranking Mechanisms,” *American Economic Review*, 102, 325–329.
- CARROLL, G. (2014): “A Complexity Result for Undominated-Strategy Implementation,” *Working Paper*.
- (2019): “Robustness in Mechanism Design and Contracting,” *Annual Review of Economics*, 11, 139–166.
- CHEN, Y.-C. AND J. LI (2018): “Revisiting the Foundations of Dominant-Strategy Mechanisms,” *Journal of Economic Theory*, 178, 294–317.
- CHERNEV, A., U. BÖCKENHOLT, AND J. GOODMAN (2015): “Choice Overload: A Conceptual Review and Meta-Analysis,” *Journal of Consumer Psychology*, 25, 333–358.
- CHUNG, K.-S. AND J. C. ELY (2007): “Foundations of Dominant-Strategy Mechanisms,” *Review of Economic Studies*, 74, 447–476.
- CRAMTON, P., R. GIBBONS, AND P. KLEMPERER (1987): “Dissolving a Partnership Efficiently,” *Econometrica*, 55, 615–632.
- GLAZER, J. AND A. RUBINSTEIN (2012): “A Model of Persuasion with Boundedly Rational Agents,” *Journal of Political Economy*, 120, 1057–1082.
- (2014): “Complex Questionnaires,” *Econometrica*, 82, 1529–1541.
- HAGERTY, K. M. AND W. P. ROGERSON (1987): “Robust Trading Mechanisms,” *Journal of Economic Theory*, 42, 94–107.

- HARTLINE, J. D. AND T. ROUGHGARDEN (2009): “Simple versus Optimal Mechanisms,” in *Proceedings of the 10th ACM conference on Electronic commerce*, 225–234.
- JACKSON, M. O. (1992): “Implementation in Undominated Strategies: A Look at Bounded Mechanisms,” *Review of Economic Studies*, 59, 757–775.
- JAKOBSEN, A. M. (2020): “A Model of Complex Contracts,” *American Economic Review*, 110, 1243–1273.
- JEHIEL, P., M. MEYER-TER VEHN, B. MOLDOVANU, AND W. R. ZAME (2006): “The Limits of Ex Post Implementation,” *Econometrica*, 74, 585–610.
- LI, S. (2017): “Obviously Strategy-Proof Mechanisms,” *American Economic Review*, 107, 3257–87.
- LOERTSCHER, S. AND L. M. MARX (2020): “A Dominant-Strategy Asset Market Mechanism,” *Games and Economic Behavior*, 120, 1–15.
- LU, H. AND J. ROBERT (2001): “Optimal Trading Mechanisms with Ex Ante Unidentified Traders,” *Journal of Economic Theory*, 97, 50–80.
- MACKENZIE, A. (2019): “A Revelation Principle for Obviously Strategy-Proof Implementation,” *Working Paper*.
- MAILATH, G. AND A. POSTLEWAITE (1990): “Asymmetric Information Bargaining Problems with Many Agents,” *Review of Economic Studies*, 57, 351–367.
- MALAKHOV, A. AND R. V. VOHRA (2009): “An Optimal Auction for Capacity Constrained Bidders: A Network Perspective,” *Economic Theory*, 39, 113–128.
- MUKHERJEE, S., N. MUTO, E. RAMAEKERS, AND A. SEN (2019): “Implementation in Undominated Strategies by Bounded Mechanisms: The Pareto Correspondence and a Generalization,” *Journal of Economic Theory*, 180, 229–243.
- MUKHERJEE, S., N. MUTO, AND A. SEN (2020): “Implementation in Undominated Strategies with Applications to Economic Environments,” *Working Paper*.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- MYERSON, R. AND M. SATTERTHWAITE (1983): “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 29, 265–281.
- NISAN, N. AND A. RONEN (2001): “Algorithmic Mechanism Design,” *Games and Economic behavior*, 35, 166–196.
- PYCIA, M. AND P. TROYAN (2021): “A Theory of Simplicity in Games and Mechanism Design,” *Working Paper*.
- RAY, D. & A. ROBSON (2018): “Certified Random: A New Order for Coauthorship,” *American Economic Review*, 108, 489–520.
- ROCHET, J.-C. (1987): “A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context,” *Journal of Mathematical Economics*, 16, 191–200.

- SPIEGLER, R. (2011): *Bounded Rationality and Industrial Organization*, Oxford University Press.
- TILLIO, A. D., N. KOS, AND M. MESSNER (2017): “The design of ambiguous mechanisms,” *Review of Economic Studies*, 84, 237–276.
- YAMASHITA, T. (2015a): “Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism,” *Review of Economic Studies*, 82, 1223–1246.
- (2015b): “Strategic and structural uncertainty in robust implementation,” *Journal of Economic Theory*, 159, 267–279.
- YAMASHITA, T. AND S. ZHU (2020): “On the Foundations of Ex Post Incentive Compatible Mechanisms,” *American Economic Journal: Microeconomics*, Forthcoming.

A Appendix

A.1 The optimal OSP mechanism for Example 1

We derive the optimal OSP mechanism in Example 1. It follows from Theorem 3.1 of [Bade and Gonczarowski \(2017\)](#) that it is without loss of optimality for OSP implementation to look at gradual revelation mechanisms (see [Ashlagi and Gonczarowski \(2018\)](#), [Pycia and Troyan \(2021\)](#) and [Mackenzie \(2019\)](#) for related revelation principles for OSP mechanisms). In our simple example with two players and two types, this means that we can assume that in the best OSP mechanism, at the first decision node, one of the players (“leader”) makes a binary decision (with the two types choosing different actions—potentially leading to the same outcome—as part of their obviously dominant strategies), and then in each of the two possible histories, having observed the choice of the leader, the other player (“follower”) makes a binary decision.

Therefore, an upper bound on the profit in the optimal OSP mechanism can be derived in the following way. When the follower chooses her optimal action, she already knows the action chosen by the leader. OSP requires that for any choice of the leader, each type of the follower must weakly prefer her equilibrium strategy (action) to choosing the alternative action. Thus, each type of the follower must have a standard *dominant strategy* in the normal-form representation of the game. In contrast, when the leader chooses her action at the initial decision node, it must be that the worst possible payoff from choosing her equilibrium action (over the two possible actions that can be selected by the follower in the subgame) is weakly higher than the best possible payoff from choosing the alternative action. In the normal-form representation of the game, this can be captured by requiring that the payoff from the equilibrium strategy of each type of

the leader under any choice of the strategy S_f for the follower is weakly higher than the payoff from following the alternative strategy under any choice of the strategy S'_f (where, importantly, S_f could be different from S'_f).

Summarizing, an upper bound can be derived by using a normal-form game with the usual strategy-proof constraints, except that for the leader the constraints are strengthened in the way described above. Crucially, all these constraints are linear in the allocation and transfers, and so is the objective function of the designer (intuitively, we avoid taking the min and max in the definition of obvious dominance by iterating over all possible *pairs* of strategies (S_f, S'_f) for the follower when comparing the two strategies of the leader). Thus, we obtain a linear program that can be numerically solved using a standard linear programming solver, yielding an upper bound of $1/5$. This upper bound is achieved by the OSP mechanism described in the example, proving its optimality.

A.2 Proof of Proposition 1

Fix a solution concept K . If the AAT property fails in some implementation environment, then there exists a type θ_i of agent i and an outcome function $\lambda \in \Lambda(\Theta \setminus \{\theta_i\})$ such that for all $\bar{\lambda} \in \Lambda(\Theta)$, $\bar{\lambda}|_{\Theta \setminus \{\theta_i\}} \neq \lambda$. We can assume without loss of generality that there exists such λ that is deterministic for any θ , since if the AAT property holds for all deterministic outcome functions, then it also must hold for all stochastic outcome functions. Define an objective function v of the designer by

$$v(x, \theta) = \mathbf{1}_{\{x=\lambda(\theta), \theta \in \Theta \setminus \{\theta_i\}\}}.$$

Clearly, the best simple mechanism on the type space $\Theta \setminus \{\theta_i\}$ yields an expected payoff of 1 to the designer.

Fix an arbitrary full-support distribution of types π . With the objective function v on Θ , the best simple mechanism on the type space Θ must yield an optimal payoff strictly lower than $1 - p$, where p is the unconditional probability of type θ_i under π , since it is not possible, by assumption, to implement the outcome function λ on $\Theta \setminus \{\theta_i\}$. We denote the optimal payoff by $b < 1 - p$.

We will show that there exists a complex mechanism that generates a strictly higher expected payoff on Θ than b , regardless of how agents behave when they are confused. Let Γ be a mechanism that implements the outcome function λ on $\Theta \setminus \{\theta_i\}$. Consider the same mechanism Γ on Θ . By assumption, Γ is not simple on Θ because type θ_i is strategically confused. But because it is simple on $\Theta \setminus \{\theta_i\}$, all types in $\Theta \setminus \{\theta_i\}$ still play the same (K -dominant) strategy. Therefore, no matter how type θ_i behaves,

the designer gets an expected payoff of at least $1 - p$ times 1, which is strictly higher than b , thereby strongly dominating the optimal simple mechanism.

Now suppose that the AAT property fails in some design environment. Then, there exists some distribution π' supported on $\Theta \setminus \{\theta_i\}$ such that the best simple mechanism Γ on $\Theta \setminus \{\theta_i\}$ gives an expected payoff b to the designer, while the best simple mechanism on Θ (that is, with incentive constraints also imposed for type θ_i) gives some smaller expected payoff $b' < b$. Specify π on Θ to coincide with π' conditional on θ_i not realizing, and let θ_i have prior probability $\epsilon > 0$ under π (it does not matter for our argument how the types of remaining players are distributed conditional on θ_i). Then, the best simple mechanism under π can yield a payoff of at most $(1 - \epsilon)b' + \epsilon\bar{v}$, where \bar{v} is an upper bound on the designer's payoff when player i has type θ_i . If the designer instead uses the same mechanism Γ that yields b on $\Theta \setminus \{\theta_i\}$, then her expected payoff is at least $(1 - \epsilon)b + \epsilon\underline{v}$, where \underline{v} is a lower bound on the designer's payoff when player i has type θ_i . Thus, when ϵ is small enough, Γ (complex on Θ) strongly dominates the best simple mechanism on Θ .

A.3 Proof of Claim 1

We will explicitly construct a K -dominant strategy for the additional type θ (we drop the subscript, since there is only one player) taking as given any simple mechanism Γ on $\Theta \setminus \{\theta\}$. Under our definitions of all three solution concepts, it is immediate (see [Ashlagi and Gonczarowski \(2018\)](#) and [Pycia and Troyan \(2021\)](#) for formal arguments) that we can obtain a mechanism Γ' that implements the same outcome function as Γ by letting Nature move only once in the first node of the extensive form game. Moreover, the move by Nature is observed, and the agent has a single decision node at which she effectively selects an outcome x from some menu of choices $\mathcal{M} \subset \mathcal{X}$, where the menu \mathcal{M} may depend on Nature's choice. For any realized menu \mathcal{M} , specify that type θ chooses her most preferred outcome from that menu. Clearly, this strategy is K -dominant for type θ . At the same time, all remaining types preserve their K -dominant strategies since the extensive-form game Γ' has not been modified.

A.4 Proof of Claim 2

We explicitly construct an example in which the best simple mechanism is strongly dominated for all $K \in \{\text{SP}, \text{OSP}, \text{SOSP}\}$. There is one player with three equally-likely types, $\Theta = \{u, m, d\}$, $\mathcal{X} = \{U, U', M, M', D, D'\}$, with the following preferences:

1. type u : $M > U > D' > D > M' > U'$;
2. type m : $D > M > U' > U > D' > M'$;

3. type d : $U > D > M' > M > U' > D'$.

The designer gets a utility of 1 if the type is $j \in \{u, m, d\}$ and she implements outcome J or J' ; she gets -1 otherwise.

Lemma 1. *The best simple mechanism is to implement any fixed (possibly random) outcome; the expected payoff for the designer is $-1/3$.*

It suffices to argue that this is true with SP as the solution concept. Recall that it is without loss of optimality for the designer to consider deterministic mechanisms when optimizing in the class of simple mechanisms. A deterministic SP mechanism for a single agent can be represented as a direct assignment of alternatives to types such that no type strictly prefers another type's assignment to her own. Suppose that there exists such an assignment that gives the designer an expected payoff strictly above $-1/3$. Then, at least two types $j \in \{u, m, d\}$ must be assigned either J or J' . If any type j is assigned J' , then no other alternative can be offered by the mechanism since j ranks J' last. Thus, the mechanism must offer J to two distinct types j . However, that's a contradiction because at least one of these types would prefer the allocation of the other one, no matter which two types $j \in \{u, m, d\}$ we choose.

To finish the proof, we construct a superior complex mechanism Γ .

Lemma 2. *There exists a complex mechanism Γ that guarantees the designer an expected payoff of 0.*

In the mechanism Γ , the designer uses a randomization device that has equally likely outcomes H and T , and offers three possible strategies to the agent, as represented by the following normal form:

	H	T
S_U	U	D'
S_M	M	U'
S_D	D	M'

For any solution concept $K \in \{\text{SP}, \text{OSP}, \text{SOSP}\}$, type $j \in \{u, m, d\}$ is confused between the two strategies offering J and J' , respectively, with the former one (denoted S_J) leading to the 2nd or 3rd alternative, and the latter to the 1st or 6th alternative. However, the remaining strategy is K -dominated for j by the strategy S_J as it leads to the 4th or 5th alternative. Regardless of how j resolves her strategic confusion, either J or J' is implemented with probability $1/2$, and thus the designer obtains 0 in expectation.

A.5 Proof of Proposition 2

We first show that if f is an ordinal incentive-compatible allocation where the coefficients $f(\theta_i)(x_j)$ for all i and j are rational numbers, then f can be implemented via a complex mechanism. When all the coefficients are rational numbers, there exists some integer \bar{N} such that $f(\theta_i)(x_j) \cdot \bar{N}$ is an integer for all i and j . We explicitly construct a mechanism that involves a randomization device with \bar{N} outcomes, all with equal probability of $1/\bar{N}$. For each $\theta \in \Theta$, let S_θ be the set of all possible mappings $d : \{1, 2, \dots, \bar{N}\} \rightarrow \mathcal{X}$ with the property that $|k : d(k) = x_j| = f(\theta)(x_j) \cdot \bar{N}$ for all j (the choice of \bar{N} guarantees that there are exactly $N!$ such mappings for each θ). The mechanism is static: The agent chooses one of the mappings from one of $\{S_\theta\}_{\theta \in \Theta}$, and then the realization of the randomization device determines the final outcome. Since f is ordinal incentive-compatible, for any $s' \notin S_\theta$, by construction, there exists a strategy $s \in S_\theta$ such that either (i) s and s' are duplicate strategies or (ii) s SP-dominates s' for θ . Therefore, regardless of how type θ resolves her uncertainty, she chooses a strategy from S_θ , and hence the outcome implemented conditional on θ has the distribution $f(\theta)$. Therefore, this complex mechanism implements f .

Now suppose that there exists an ordinal incentive-compatible allocation that yields a strictly higher expected payoff for the designer than the best deterministic mechanism. Since the designer is an expected utility maximizer, there exists an ordinal incentive-compatible allocation where all the coefficients are rational numbers that yields a strictly higher expected payoff for the designer than the best deterministic mechanism. This, coupled with the arguments in the previous paragraph and the observation that there always exists an optimal simple mechanism that is deterministic, establishes the desired result.

A.6 Proof of Proposition 3 and 3'

We prove Proposition 3 first, and then explain how to modify the steps to obtain Proposition 3'. We start with a lemma that builds on the insight in Yamashita (2015a, Theorem 1), and can be viewed as its generalization to a larger class of environments and the solution concepts of OSP and SOSP.

Lemma 3. *For any mechanism Γ , $\min V(\Gamma)$ is upper-bounded by the value of the IC-relaxed problem corresponding to any fixed collection of trees $\{T_i\}_{i \in \mathcal{N}}$ (as defined in Section 3.3).*

Proof. Fix an arbitrary finite mechanism Γ , an agent i , and a tree T_i . Let $T_i^+(\theta_i) = \{\theta'_i : \theta'_i \rightarrow \theta_i\}$ be the set of types who point towards type θ_i in the tree T_i . Consider the

following procedure. Starting at the root of the tree T_i —which is some type θ_i^0 with no edges coming out of it—select any K -undominated strategy for θ_i^0 , $S_i^0 \in U_i^K(\theta_i^0)$. Next, for any type $\theta_i' \in T_i^+(\theta_i^0)$, we can find an undominated strategy $S_i' \in U_i^K(\theta_i')$ that either K -dominates S_i^0 or is equal to S_i^0 (this step uses finiteness of the mechanism; if S_i^0 is not in $U_i^K(\theta_i')$, then there must exist a strategy in $U_i^K(\theta_i')$ that K -dominates it). We proceed inductively. Once some type θ_i is assigned a strategy, we assign undominated strategies to all types $T_i^+(\theta_i)$ that either equal or K -dominate the strategy assigned to θ_i . Because T_i is a finite tree, this procedure must stop at some point, with every type being assigned a strategy. The same procedure is carried out for all other agents.

Because each type is assigned a strategy from $U_i^K(\cdot)$ in the procedure, when all types execute their assigned strategies, the expected payoff \bar{v} to the designer must weakly exceed $\min V(\Gamma)$ (which is the outcome of the designer-adversarial selection from $U_i^K(\cdot)$). Moreover, the procedure guarantees that the mechanism—along with the assignment of strategies—is feasible for the IC-relaxed problem corresponding to the collection of trees $\{T_i\}_{i \in \mathcal{N}}$. Therefore, \bar{v} , and hence also $\min V(\Gamma)$, is weakly below the value of the IC-relaxed problem. \square

Proposition 3 follows immediately: By assumption, there exists a collection of trees $\{T_i\}_{i \in \mathcal{N}}$ such that the value of the IC-relaxed problem corresponding to $\{T_i\}_{i \in \mathcal{N}}$ is the same as the designer’s expected payoff from the optimal simple mechanism. Hence, by applying Lemma 3 for the collection $\{T_i\}_{i \in \mathcal{N}}$, we conclude that there cannot exist a mechanism Γ with $\min V(\Gamma)$ strictly above the expected payoff of the optimal simple mechanism, and hence the optimal simple mechanism is not strongly dominated.

We now explain how to incorporate participation constraints in the above procedure to obtain Proposition 3'. Suppose that the superior complex mechanism is required to provide partial incentives to participate. In that case, in the proof of Lemma 3, we can select an undominated strategy for type θ_i^0 —the root of the tree T_i —that satisfies the IR constraint (1) (such a strategy exists by the definition of partial incentives to participate). Thus, we can obtain a version of Lemma 3 that assumes that Γ provides partial incentives to participate and concludes that $\min V(\Gamma)$ is upper-bounded by the value of the IC-relaxed and IR-relaxed problem. If the superior complex mechanism is required to provide full incentives to participate, then we know that all undominated strategies must satisfy (1). Thus, we can obtain a version of Lemma 3 which assumes that Γ provides full incentives to participate and concludes that $\min V(\Gamma)$ is upper-bounded by the value of the IC-relaxed and IR-non-relaxed problem.

A.7 Proof of Proposition 4

For each solution concept $K \in \{\text{SP}, \text{OSP}, \text{SOSP}\}$, we show that the simple mechanism is weakly dominated by explicitly constructing a complex mechanism that dominates it. The mechanism we construct can be interpreted as a delegation mechanism in which one agent is delegated to choose a simple mechanism (for the agents to play) from a menu of two simple mechanisms specified by the designer.²⁹ We present the proof for each solution concept separately.

SP: Fix a player i . We add a new node for player i from which play begins. Player i chooses from two options: “Choose Γ ” or “Choose $\Gamma_{-i}^{\mathcal{Y}}$.” If she chooses Γ , the game tree is the one associated with Γ . If she chooses $\Gamma_{-i}^{\mathcal{Y}}$, the game tree is the one associated with the game $\Gamma_{-i}^{\mathcal{Y}}$. We call this new composite game Γ' .

First, we claim that all players other than i have a dominant strategy in Γ' . This is immediate from the fact that each player $-i$ has a dominant strategy in Γ and a dominant strategy in $\Gamma_{-i}^{\mathcal{Y}}$. Second, all types of player i not in $\Theta_i^{\mathcal{Y}}$ also have a dominant strategy which is to choose Γ , and then follow the same strategy $\mathcal{S}_i(\theta_i)$ that was dominant for Γ . This is immediate from the fact that for these types, the best that can happen after choosing $\Gamma_{-i}^{\mathcal{Y}}$ cannot be better than the worst possible outcome in Γ when they follow $\mathcal{S}_i(\theta_i)$. Third, we claim that for all types $\theta_i \in \Theta_i^{\mathcal{Y}}$, the option to choose $\Gamma_{-i}^{\mathcal{Y}}$ is SP-undominated. Indeed, fix the strategy profile \bar{S}_{-i} that yields the minimum in the definition of $\Theta_i^{\mathcal{Y}}$, and—using condition 1 in the proposition—let $\bar{S}_{-i}^{\mathcal{Y}}$ be the profile that leads to the outcome $x^* \in \operatorname{argmax}_{x \in \mathcal{Y}} u_i(x, \theta_i)$ in the game $\Gamma_{-i}^{\mathcal{Y}}$. Then, if players $-i$ follow the strategy $(\bar{S}_{-i}, \bar{S}_{-i}^{\mathcal{Y}})$ in Γ' , by definition of $\Theta_i^{\mathcal{Y}}$, the best response for type θ_i is to choose the game $\Gamma_{-i}^{\mathcal{Y}}$ at the first decision node.

Overall, it is an undominated strategy for types $\theta_i \in \Theta_i^{\mathcal{Y}}$ to choose the game $\Gamma_{-i}^{\mathcal{Y}}$, and when they do, play among players $-i$ in that subgame proceeds as in the original game $\Gamma_{-i}^{\mathcal{Y}}$. By condition 2 in the proposition, the designer receives a higher (sometimes strictly) conditional expected payoff in that case, compared to the conditional expected payoff she would have received in the game Γ . Therefore, the game Γ is weakly dominated by Γ' .

OSP: The game Γ' is defined just as in the SP case. It suffices to prove that the three claims made in the second paragraph for the case of SP dominance continue to hold for OSP dominance. The first claim extends trivially. As for the second claim, $\theta_i \notin \Theta_i^{\mathcal{Y}}$ is equivalent to saying that choosing the game $\Gamma_{-i}^{\mathcal{Y}}$ is obviously dominated by the strategy

²⁹These complex mechanisms are “type 1 strategically simple” according to the notion in [Börger and Li \(2019\)](#).

“choose Γ ” and then play according to $\mathbf{S}_i(\theta_i)$, since the unique point of departure for these two strategies is the first node for player i . Finally, the last claim follows from the fact that if a strategy is not SP-dominated for a type, then it is also not OSP-dominated.

SOSP: For this case, we modify the construction of the game Γ' . [Pycia and Troyan \(2021\)](#) prove that it is without loss of generality for the designer to choose a game form in which each player only moves once in every history, in the sense that for any SOSP mechanism, there exists a payoff equivalent SOSP mechanism (for all players and the designer) with that property. Thus, we can replace both Γ and $\Gamma_{-i}^{\mathcal{Y}}$ with their payoff-equivalent versions in which each player moves at most once in every history. Moreover, [Pycia and Troyan \(2021\)](#) show that, without loss of generality, Nature’s moves take place at the beginning of each history, if at all.

Let i be the player that moves first in the game Γ (if Nature moves first, i can depend on Nature’s choice, and we apply the argument case-by-case, for every realization of Nature’s choice). For that player, add a new potential action from the initial node which is to choose the game $\Gamma_{-i}^{\mathcal{Y}}$. If that action is taken by player i , play proceeds to the game tree $\Gamma_{-i}^{\mathcal{Y}}$. This defines a new game Γ' .

Note that in Γ' all players still move at most once in every history; player i has no moves in the game $\Gamma_{-i}^{\mathcal{Y}}$, while players $-i$ either play the game Γ (in which they have at most one move) or the game $\Gamma_{-i}^{\mathcal{Y}}$ (in which they have at most one move). Therefore, all the conclusions from the OSP case exceed to this case (since OSP and SOSP coincide if a player only moves once).

A.8 Proof of Claim 3

We first consider the construction for the case $N = 2$ without relying on randomization. We then cover the case $N = 1$.

Fix a simple mechanism Γ , with some expected revenue R . Suppose that $N \geq 2$. For any i , define $\mathcal{Y} \subset \mathcal{X}$ to contain two outcomes: (1) Player i pays M to the designer while some player j receives $\epsilon > 0$ from the designer, and (2) Player i receives M from the designer while player j receives 0 from the designer. The mechanism $\Gamma_{-i}^{\mathcal{Y}}$ is defined as having just one information node for player j who chooses between the two options in \mathcal{Y} . This satisfies condition 1 of Proposition 4. Moreover, the mechanism $\Gamma_{-i}^{\mathcal{Y}}$ is simple because player j has a K -dominant strategy to select option (1)—player i pays M to the designer while player j receives $\epsilon > 0$ from the designer. When M is large enough, $\Theta_i^{\mathcal{Y}} = \Theta_i$, and when additionally ϵ is small enough, condition 2 of Proposition 4 holds, since the conditional expected payoff from the mechanism $\Gamma_{-i}^{\mathcal{Y}}$ is unbounded in M . Thus,

by Proposition 4, the simple mechanism Γ is weakly dominated. It remains to check that the weakly-dominating complex mechanism provides partial incentives to participate. This is immediate from the proof of Proposition 4 that explicitly constructs the weakly-dominating complex mechanism Γ' : Intuitively, in Γ' , player i chooses between the games Γ and $\Gamma_{-i}^{\mathcal{Y}}$; thus, since each type of player i had a K -dominant strategy satisfying (1) in Γ , each type continues to have at least one K -undominated strategy satisfying (1) in Γ' .

Consider the case $N = 1$. Let $\mathcal{Y} \subset \mathcal{X}$ contain two outcomes: (1) Player i pays M to the designer, and (2) Player i receives M from the designer. The mechanism $\Gamma_{-i}^{\mathcal{Y}}$ is defined as having just one information node for Nature that chooses option (1) with probability $1 - \epsilon$ and option (2) with probability ϵ . For ϵ sufficiently small, by Proposition 4 and the same arguments as before, the mechanism Γ is weakly dominated. (Of course, this construction could also be used for general N .)

A.9 Proof of Claim 4

Fix Γ that is a revenue-maximizing simple mechanism. Because of symmetry, we can assume without loss of optimality for the designer that Γ is symmetric as well, so we only have to check the assumptions of Proposition 4 for some player i . Define \mathcal{Y} to contain two options: (1) Player i wins the object and pays \bar{u} while some player j receives $\epsilon > 0$, and (2) Player i wins the object and pays $\bar{u} - \delta$ while player j receives 0. In the game $\Gamma_{-i}^{\mathcal{Y}}$, player j selects one option from \mathcal{Y} . $\Gamma_{-i}^{\mathcal{Y}}$ is thus simple, with player j always selecting option (1). Take $\delta > 0$ small enough so that $\Theta_i \cap (\bar{u} - \delta, \bar{u}) = \emptyset$ (using finiteness of the type space). Then, we have $\Theta_i^{\mathcal{Y}} = \{\bar{u}\}$, because for all other types the payoff from \mathcal{Y} is strictly negative (while the original strategy in Γ guarantees a non-negative payoff).³⁰ Moreover, type \bar{u} always gets a non-negative payoff in $\Gamma_{-i}^{\mathcal{Y}}$ so she has a full incentive to participate. Condition 1 of Proposition 4 holds, while condition 2 is satisfied as long as the mechanism Γ extracts less than \bar{u} from type \bar{u} of player i . Thus, by Proposition 4, Γ is weakly dominated as long as it is not a mechanism that offers a price \bar{u} to player i . By symmetry, the only case not covered by the argument is when Γ is payoff-equivalent to a mechanism that offers a price \bar{u} to all players in random order. However, such a mechanism is suboptimal by assumption.

A.10 Proof of Claim 5

By assumption, the unique optimal simple mechanism is outcome-equivalent to offering some price p^* ; let \bar{v} denote the corresponding optimal expected revenue. Note that, by

³⁰ $\bar{u} \in \Theta_i^{\mathcal{Y}}$ as long as there is positive probability that she does not receive the object in Γ —that is guaranteed by the assumption that the players and Γ are symmetric.

optimality, there must exist a type $\theta = p^*$ in Θ .

Towards a contradiction, suppose that there exists a weakly dominating complex mechanism Γ . By definition, $\min V(\Gamma) \geq \bar{v}$; we will first show that $\min V(\Gamma) = \bar{v}$. We claim that there exists a selection from the set of K -undominated strategies $U_i^K(\cdot)$ in Γ such that (i) if every type plays the assigned strategy, the expected revenue for the designer is equal to at least $\min V(\Gamma)$, and (ii) local downward incentive constraints hold, that is, the strategy assigned to type θ_i K -dominates (for θ_i) the strategy assigned to the highest type lower than θ_i . This follows from Lemma 3 in Appendix A.6 that establishes a more general result. Moreover, it is well known that for revenue maximization with a single player under SP and OSP, only local downward incentive constraints bind in the optimal mechanism; hence, the expected revenue under the selection of undominated strategies described above is upper-bounded by \bar{v} . Since $\bar{v} \leq \min V(\Gamma)$, we conclude that $\min V(\Gamma) = \bar{v}$ and there exists a selection from the set of undominated strategies which satisfies local downward incentive constraints and yields an expected revenue of \bar{v} .

By the assumption that the optimal simple mechanism is unique, the only way to generate the expected revenue of \bar{v} while satisfying local downward incentive constraints is for all types weakly above p^* to buy for sure at the expected price of p^* . By the full incentive to participate, type p^* must have a strategy under which she buys for sure at a price of p^* in every history. In particular, Γ must offer such a strategy. Moreover, again by the full incentive to participate, this strategy must be K -dominant for type p^* .

By the assumption that Γ weakly dominates the simple mechanism, there must exist a strategy S_1 that is K -undominated for some type θ , and a strategy for Nature S_0 such that the agent pays $q > p^*$ in the outcome $g(z(h_\theta, S_0, S_1))$. Because the strategy “buy for sure at a price of p^* ” is available, for S_1 to be K -undominated, it must sometimes (for some strategy of Nature) generate an outcome “buy with probability x at a price of r ” that is strictly preferred by θ to the outcome “buy for sure at a price of p^* .” At the same time, type p^* cannot derive positive utility from the outcome “buy with probability x at a price of r ” as otherwise the strategy S_1 would be undominated for type p^* , violating the full incentive to participate (since S_1 sometimes leads to the agent paying $q > p^*$). Moreover, note that $\theta > p^*$, by the full incentive to participate. We are ready to obtain a contradiction: By the above reasoning, we have $\theta x - r > \theta - p^*$ and $p^*x - r \leq 0$. This implies

$$\theta - p^* < \theta x - r \leq \theta x - p^*x = x(\theta - p^*) \leq \theta - p^*,$$

a contradiction.

Supplemental Material

OA.1 Supplemental Material for Subsection 3.3

In this appendix, we give a formal definition of the uniform shortest-path tree condition. We work in a quasi-linear setting as defined in Subsection 4.2, with X being the set of all (possibly random) physical allocations, and $\mathcal{X} = X \times \mathbb{R}^N$. We say that a decision rule $g : \Theta \rightarrow X$ is strategy-proof if there exists a transfer scheme $t = (t_1, \dots, t_N)$ such that (g, t) is an SP mechanism (treated as a direct mechanism). The following standard lemma is due to [Rochet \(1987\)](#).

Lemma 4. *A necessary and sufficient condition for a decision rule g to be strategy-proof is the following cyclical monotonicity condition: $\forall i \in \mathcal{N}, \forall \theta_{-i} \in \Theta_{-i}$, and for every sequence of types $(\theta_{i,1}, \theta_{i,2}, \dots, \theta_{i,m})$ with $\theta_{i,m} = \theta_{i,1}$, we have*

$$\sum_{n=1}^{m-1} \left[u_i(g(\theta_{i,n}, \theta_{-i}), \theta_{i,n+1}) - u_i(g(\theta_{i,n}, \theta_{-i}), \theta_{i,n}) \right] \leq 0. \quad (\text{OA.2})$$

We first collect some graph-theoretic terminology in Definition 10. Instead of imposing an IR constraint (1) for the root of each tree, we introduce a “dummy” type θ_0 (which will be the new root) that corresponds to not participating in the mechanism.

Definition 10. *Fix a strategy-proof decision rule g , agent $i \in \mathcal{N}$, and θ_{-i} .*

- (1) *The set of nodes is $\Theta_i \cup \{\theta_0\}$;*
- (2) *For any $\theta_i \in \Theta_i$ and $\theta'_i \in \{\Theta_i \setminus \{\theta_i\}\} \cup \{\theta_0\}$, $\theta_i \rightarrow \theta'_i$ is a directed edge with length*

$$u_i(g(\theta_i, \theta_{-i}), \theta_i) - u_i(g(\theta'_i, \theta_{-i}), \theta_i);$$

- (3) *A path from the dummy type θ_0 to $\theta_{i,m} \in \Theta_i$ is a sequence of nodes $P = (\theta_0, \theta_{i,1}, \theta_{i,2}, \dots, \theta_{i,m})$ where (i) $\theta_{i,j} \in \Theta_i, \forall j = 1, 2, \dots, m$; and (ii) $j \neq j' \implies \theta_{i,j} \neq \theta_{i,j'}$.*

Definition 11. *Fix a strategy-proof decision rule g , agent $i \in \mathcal{N}$, and θ_{-i} . A shortest-path tree is the union of shortest paths from the root to all nodes such that if θ'_i belongs to the shortest path from the root θ_0 to some $\theta_i \in \Theta_i$, then the truncation of the path from θ_0 to θ'_i defines the shortest path from θ_0 to θ'_i .*

Definition 12. We say that the uniform shortest-path tree condition holds if, for each agent $i \in \mathcal{N}$, there is the same shortest-path tree for all strategy-proof decision rules g and other agents' reports θ_{-i} .

When the uniform shortest-path tree condition is satisfied, we can drop the dependence of the shortest-path tree on g and θ_{-i} . Let E_s denote the collection of the directed edges in the uniform shortest-path tree.

Suppose that the uniform shortest-path tree condition holds. In the optimal SP mechanism design problem, it suffices to consider constraints that correspond to the edges on the uniform shortest-path tree, subject to the decision rule g satisfying the cyclical monotonicity constraint (OA.2). As is standard in the mechanism design literature, we first consider the following relaxed problem, in which we ignore the cyclical monotonicity constraint. A regularity condition on π is then imposed to ensure that some optimal decision rule g^* that solves the relaxed maximization problem automatically satisfies the cyclical monotonicity constraint.

$$\begin{aligned} & \max_{g(\cdot) \in X, t_i(\cdot)} \sum_{\theta \in \Theta} \pi(\theta) \sum_{i \in \mathcal{N}} t_i(\theta) && \text{(SP-relaxed)} \\ & \text{subject to } \forall i \in \mathcal{N}, \forall (\theta_i, \theta'_i) \in E_s, \forall \theta_{-i} \in \Theta_{-i}, \\ & u_i(g(\theta_i, \theta_{-i}), \theta_i) - t_i(\theta_i, \theta_{-i}) \geq u_i(g(\theta'_i, \theta_{-i}), \theta_i) - t_i(\theta'_i, \theta_{-i}). \end{aligned}$$

Definition 13. We say that π is regular if the cyclical monotonicity constraint (OA.2) is automatically satisfied for some g^* that solves the optimization problem (SP-relaxed).³¹

OA.2 Supplemental Material for Subsection 4.2

In this appendix, we show that the uniqueness assumption in Claim 5 is needed. We construct an example in which the optimal SP and OSP mechanisms are not unique, and they are weakly dominated.

The agent has value 1 or 2, with equal probabilities. The optimal SP and OSP mechanisms generate an expected revenue of 1 (which can be obtained by charging a price of 1, or a price of 2). The weakly dominating mechanism features Nature that plays H with some small probability $\epsilon > 0$, and T otherwise. The mechanism offers three strategies to the agent, where the first number in every cell is the probability of trade, and the second number is the payment to the designer.

³¹To the best of our knowledge, there is no formal definition of regularity in general environments. Our definition of regularity captures how it has been used in the literature; see, for example, Myerson (1981) and Chung and Ely (2007). See Chen and Li (2018) for the primitive condition for the regularity condition in a variety of settings.

	H	T
S_1	$(1, 3/2)$	$(1, 3/2)$
S_2	$(1, 1)$	$(1, 2)$
S_3	$(1/2, 1/2)$	$(1/2, 1/2)$

For type 2, S_1 and S_2 are undominated (playing S_3 can be ruled out under the assumption that the designer can choose between strategies that are payoff equivalent, in this case S_1 and S_3). Type 1 has a dominant strategy S_3 . Thus, Γ provides full incentives to participate. In the worst case, type 2 plays S_1 , and the designer obtains an expected revenue of 1. In the best case, type 2 plays S_2 , and the designer obtains an expected revenue of $1/2 + 1/2 \cdot (2 - \epsilon)$ which can be arbitrarily close to $3/2$.

OA.3 Efficiency in bilateral trade

In this appendix, we consider the classical problem of designing an efficient mechanism that allows a seller of an indivisible object to trade with one potential buyer. Both agents have quasi-linear preferences and private information about their values (which are drawn from a non-degenerate distribution). It is known from [Hagerty and Rogerson \(1987\)](#) that the only SP mechanisms (and hence also OSP and SOSP mechanisms) that satisfy ex-post budget balance and individual rationality are outcome-equivalent to posted price mechanisms. In a posted price mechanism, the designer chooses a (possibly random) price, without taking into account any of the agents' private information. Trade comes about only when both agents agree to trade at that price.

Claim 6. *The best simple mechanism in the bilateral trade model is weakly dominated by a complex mechanism (that is ex-post budget-balanced and provides full incentives to participate).*³²

Proof. We will verify the assumptions of Proposition 4 with i being the seller (an analogous argument works for the buyer). Fixing Γ , let p be the posted price.³³ Let the set \mathcal{Y} contain two options: (1) No trade and (2) Trade at price $p - \epsilon$ for some $\epsilon > 0$. The mechanism $\Gamma_{-i}^{\mathcal{Y}}$ consists in the buyer choosing one option from \mathcal{Y} . Clearly, that mechanism is simple for the buyer. Condition 1 in Proposition 4 is trivially satisfied. $\Theta_i^{\mathcal{Y}}$ contains all types of the seller strictly below $p - \epsilon$. Finally, the second condition in Proposition 4 is satisfied as long as the buyer chooses option (2) with positive probability which is when her type is above $p - \epsilon$. Hence, as long as the type spaces for the buyer

³²[Börgers and Smith \(2012\)](#) show this result for the solution concept of SP.

³³If that price is random, then we fix p to be the realization that yields the highest conditional expected payoff for the designer—this can only increase the expected surplus in the mechanism Γ .

and the seller are not degenerate, we can find ϵ for which the simple mechanism Γ is weakly dominated. \square

The superior complex mechanism (that we formally construct based on Proposition 4) can be viewed as a “price cap mechanism.” The designer sets a price cap of p but gives the seller an additional strategic option to lower the price to some $p' < p$. This leads to strategic confusion of seller types below p' ; indeed, these types may benefit from lowering the price when the buyer’s value turns out to lie between p' and p . At the same time, no matter how these types behave, total surplus can only improve. Thus, this mechanism weakly dominates the posted price mechanism with price p .

OA.4 Inequity aversion in voting

In this appendix, we consider an environment without transfers. There are two agents and three alternatives: $\mathcal{X} = \{a, b, c\}$. Each agent’s type can be represented as a ranking of the three alternatives. The distribution of types π is i.i.d. uniform. The designer is inequity averse in the following sense: The designer’s payoff is 0 if one agent gets her best option and the other agent gets her worst option; it is 1 otherwise. By direct calculation, among simple mechanisms, it is optimal to let the outcome be a without eliciting the preferences of the agents. For this mechanism, there are 8 type profiles (where one agent ranks a at the top and the other agent ranks a at the bottom) such that the designer obtains utility 0; thus, the designer obtains $7/9$ in expectation.

Claim 7. *The best simple mechanism in the above voting setting is weakly dominated.*³⁴

We construct the superior mechanism below. Starting from the simple mechanism of always implementing a , player i is given an additional option to let player $-i$ choose from the menu $\{b, c\}$. In this new mechanism, (1) if player i ranks a at the top, player i will not choose the additional option; there are 4 type profiles (where agent i ranks a at the top and agent $-i$ ranks a at the bottom) such that the designer obtains utility 0; (2) if agent i ranks a at the bottom, then agent i will take the additional option, in which case the designer obtains 1 regardless of agent $-i$ ’s type; (3) if agent i ’s preference is bac , then agent i will be confused as to whether to take the additional option. If she does not, then the designer obtains utility 1 regardless of agent $-i$ ’s type; if she does, then the designer obtains utility 0 if agent $-i$ ranks c at the top; (4) similarly, if agent i ’s preference is cab , then agent i will be confused as to whether to take the additional option. If she does not, then the designer obtains utility 1 regardless of agent $-i$ ’s type;

³⁴This claim is proven directly without relying on Proposition 4 but the underlying logic is analogous.

if she does, the designer obtains utility 0 if agent $-i$ ranks b at the top. Overall, in the worst-case scenario, there are 8 type profiles such that the designer obtains utility 0; if types bac or type cab do not take the additional option, the designer does strictly better than in the best simple mechanism.