Ex Post Strategy-Proofness*

Tangren Feng[†] and Qinggong Wu[‡]

December 24, 2024

Abstract

In this paper, we study ex post strategy-proofness (ESP), which extends the concept of strategy-proofness to settings with interdependent values which are more complex intrinsically. ESP requires that truthful reporting is always optimal for every agent, regardless of the strategies *and* types of other agents. We explore which social choice functions can be implemented by ESP mechanisms with transfers. For a broad class of environments, including single-good auctions, we find that implementability is precisely characterized by two conditions: 1. Monotonicity: a "higher" type leads to a "higher" alternative; and 2. Screenability: types that are "entangled" are treated identically, where entanglement is a condition that is more serious when types are densely distributed or when preference interdependence is significant. For more general settings, we show that implementability is characterized by a generalized cyclical monotonicity condition, or equivalently, by an optimality condition derived from an induced matching problem between types and alternatives. We also discuss how implementation becomes easier if agents can reason using iterated elimination of dominated strategies, and how it becomes more challenging when transfers are not available.

Keywords: Robust mechanism design, interdependent values, strategy-proofness, dominant strategy, auction design, simplicity

1 Introduction

In this paper, we investigate strategy-proof mechanism design in environments with interdependent values, where each agent's preferences depend not only on her own private information but also on the information held by others. Such settings arise naturally in many economically significant scenarios, including common value auctions, where considering others' private information can help a bidder better evaluate a good, and in jury or committee voting, where payoff-relevant information is dispersed among jurors or committee members.

---

*We thank Tilman Börgers, Yeon-Koo Che, Yi-Chun Chen, Jon Eguia, Wei He, Jiangtao Li, Bernardo Moreno, Axel Niemeyer, Mariann Ollár, Marco Ottaviani, Justus Preusser, Roland Strausz, Philipp Strack, and participants of various seminars and conferences for helpful comments. This work was supported by the European Union - Next Generation EU [MUR- Prin 2022- Prot. J53D23004150001] (Feng) and the Early Career Scheme of the Hong Kong Research Grants Council, Project 16504023 (Wu). Tangren Feng thanks the Cowles Foundation at Yale University for its hospitality while he was working on this paper.

[†] Department of Decision Sciences and IGIER, Bocconi University. Email: *tangren.feng@unibocconi.it*.

[‡] Department of Economics, Hong Kong University of Science and Technology. Email: *wqg@ust.hk*.

In mechanism design, we often aim for strategy-proofness—a property that ensures every agent has a dominant strategy. While strategy-proof mechanisms are well understood and celebrated in private value settings for their simplicity, robustness, and fairness, research on these mechanisms in interdependent value settings remains limited. Extending insights from private value settings to scenarios with interdependent values is not straightforward, as the presence of preference interdependence introduces significant complications. This paper takes a step toward addressing this gap. Specifically, we adopt the standard (partial) implementation approach to identify which social choice functions can be implemented by a strategy-proof mechanism in interdependent value settings. In doing so, our analysis also sheds new light on the limits, possibilities, and trade-offs that strategy-proof mechanism design must confront in these more complex environments.

We begin by examining what dominance entails in interdependent value settings. According to the conventional definition—typically applied in environments without preference interdependence—a strategy is considered dominant if it guarantees optimality regardless of the strategies chosen by other agents. However, extending this definition to account for preference interdependence is not straightforward. In fact, there are two distinct approaches to this extension, each depending on the timing of the evaluation of optimality.

The first approach defines dominance as guaranteed optimality in the ex post stage. Thus, a dominant strategy maximizes an agent's ex post, i.e. realized, utility for all possible strategies chosen by other agents. In other words, such a strategy is optimal regardless of the strategies *and* types of other agents. We refer to this notion of dominance as *ex post dominance*, and a mechanism in which every agent has an ex post dominant strategy is *ex post strategy-proof (ESP)*.

The other approach defines dominance as guaranteed optimality in the interim stage. Thus, a dominant strategy maximizes an agent's interim, i.e. expected, utility for all possible strategies chosen by other agents, with the expectation taken over the types of those agents. A strategy is dominant in this sense if it is optimal regardless of the strategies chosen by the other agents, conditional on one's beliefs about those agents' types. We refer to this notion of dominance as *interim dominance*, and a mechanism in which every agent has an interim dominant strategy is *interim strategy-proof (ISP)*.

In private value settings, both ESP and ISP are equivalent to conventional strategy-proofness. However, in the presence of preference interdependence, ESP is strictly stronger than ISP. This paper focuses on mechanism design that adheres to the more demanding ESP criterion, while a companion paper (Feng and Wu (2023)) studies ISP mechanism design.

ESP mechanisms possess several desirable properties. First, they are "informationally robust", because ex post dominant strategies are optimal regardless of the beliefs agents hold about each other's types. Thus, an ESP mechanism is likely to remain effective across various information structures. Second, ESP mechanisms are "strategically robust", because ex post dominant strategies are optimal regardless of the beliefs agents hold about each others' strategies. Thus, an ESP mechanism is likely to remain effective even if the agents do not accurately predict each other's strategies. Additionally, for the agents, an ESP mechanism allows them to identify their optimal strategies without concern for strategies chosen by other agents or the private information they possess. This reduces the cognitive burden on agents and minimizes the potential for strategic errors. Furthermore, the inherent simplicity of decision-making in an ESP mechanism promotes fairness, as agents who lack understanding of their circumstances or the sophistication to manipulate the mechanism are not disadvantaged compared to more sophisticated agents.

However, these desirable properties may come at a cost: the demanding nature of ESP may lead to a reduced number of mechanisms that satisfy this criterion. With this concern in mind, the primary purpose of this paper is to explore the precise scope and limitations of ESP mechanism design.

We begin with a specific class of environments that satisfy a property referred to as "one-dimensionality". This property extends an analogous concept from private value settings, where both alternatives and each agent's types can be ranked, such that a higher-ranked type exhibits a stronger preference for a higher-ranked alternative compared to a lower-ranked type. One-dimensional environments are significant because they encompass key applications such as auctions, and because in private value settings, a simple condition—where higher-ranked types lead to (weakly) higher-ranked alternatives—is known to be necessary and sufficient for implementation. We demonstrate that in interdependent value settings, once the rankings of alternatives and types are adapted to account for preference interdependence, an analogous condition, which we refer to as "monotonicity", remains necessary for ESP implementation.

However, when preference interdependence is present, monotonicity is no longer enough characterize implementability, because multiple types of the same agent may be in a situation we refer to as "entanglement", which is specific to interdependent value settings and renders them unscreenable. Essentially, these types must be treated identically in any ESP mechanism. Therefore, we have an additional necessary condition for implementability, which we refer to as the "screenability" : two types are unscreenable if they are entangled. We prove that, when combined, monotonicity and screenability are necessary and sufficient for implementability.

The seriousness of entanglement depends on two factors: the density of types and the extent of preference interdependence. When types are densely distributed or preference interdependence is significant, it becomes difficult to screen different types. If types form a continuum, even minimal preference interdependence makes them mutually unscreenable, which greatly limits implementation. Conversely, when types are relatively sparse, entanglement does not occur despite some preference interdependence, allowing any social choice function that satisfies monotonicity to be implementable.

We then examine the most general setting, allowing for environments that do not satisfy one-dimensionality. We establish that a social choice function is implementable if and only if it satisfies a condition that generalizes the well-known cyclical monotonicity. While standard cyclical monotonicity characterizes strategy-proof implementation in private value settings (Rochet (1987)), comparing the standard and generalized versions reveals the additional incentive challenges introduced by preference interdependence. Specifically, the mechanism must ensure that misreporting is unprofitable for an agent, even when she uses the most optimistic beliefs about other agents' types to assess the outcome of such a move.

Additionally, we find that implementability can be alternatively characterized by an optimality condition related to a two-sided matching problem, where one side consists of the types of an agent and the other side consists of the alternatives. This characterization extends a similar equivalence between strategy-proof implementation and optimal matching in private value settings (Dworczak and Zhang (2017)).

In the Discussion section, we explore two variations. First, we provide an example showing that if agents can reason by iteratively eliminating dominated strategies, implementation becomes easier. Second, we show that if transfers are unavailable, implementation becomes more challenging. Specifically, we find that with a moderate degree of preference interdependence, only trivial, i.e., constant, social choice functions are implementable, regardless of the type space structure. This occurs when an agent's own private information is insufficient to compare any pair of alternatives. We apply this finding to a binary voting scenario, showing that nontrivial implementation essentially requires the environment be one of private values for at least one agent. Additionally, we examine a ternary voting situation and identify sufficient conditions under which nontrivial implementation is impossible.

Literature

The literature on strategy-proof mechanism design is extensive and has a long history. However, this body of work predominantly focuses on designing mechanisms for private value settings. While some discussions on strategy-proofness in interdependent value settings exist—such as in Crémer and McLean (1985) and Williams and Radner (1988)—a comprehensive study of this topic is still lacking. Many papers instead concentrate on ex post implementation. A social choice function is ex post implementable if truthful reporting in the corresponding direct mechanism is optimal regardless of the types of other agents, provided they also report truthfully. Research on ex post implementation suggests that nontrivial implementation can be achieved under certain conditions.[1] However, implementability becomes significantly restricted when these conditions are relaxed.[2]

ESP implementation is stronger than ex post implementation because it ensures that truthful reporting remains optimal even when other agents do not report truthfully. In light of the ex post implementation literature, we find that certain conditions conducive to ex post implementation, such as one-dimensionality and single-crossing, can be suitably extended to support ESP implementation as well. However, ESP implementation is more susceptible to preference interdependence. For instance, while many significant social choice functions are known to be ex post implementable in one-dimensional environments[3]—even with a type space as dense as a continuum—they may fail to be ESP implementable if the type space lacks sufficient discreteness.

---

[1] For example, Dasgupta and Maskin (2000) requires a one-dimensional types and a single-crossing condition for efficient implementation, while Bikhchandani (2006) allows for multi-dimensional types but necessitates the absence of allocative externalities. See also Che, Kim, and Kojima (2015), Fujinaka and Miyakawa (2020), and Pourpouneh, Ramezanian, and Sen (2020) on ex post implementation without transfers.

[2] See, for instance, Jehiel, Meyer-ter Vehn, Moldovanu, and Zame (2006) in settings with transfers, or Feng, Niemeyer, and Wu (2023) and Barberà, Berga, and Moreno (2019, 2022) in settings without transfers.

[3] In auction settings, efficient social choice functions are ex post implementable when preferences satisfy appropriate single-crossing conditions; see Crémer and McLean (1985); Maskin (1992); Dasgupta and Maskin (2000); Jehiel and Moldovanu (2001); Bergemann and Välimäki (2002); Perry and Reny (2002).

In a companion paper (Feng and Wu (2023)), we analyze *interim* strategy-proofness (ISP), a weaker extension of strategy-proofness tailored for interdependent value settings. ISP relaxes the requirement for ex post optimality, and this relaxation is significant. Generally, ISP implementability imposes fewer constraints on the environment, allowing a broader range of social choice functions to be implementable as compared to ESP. For instance, in binary voting or single-good auctions, many nontrivial social choice functions are ISP implementable [4] regardless of the preference or information structure, which, as we show in this paper, is not true under ESP.

Our analysis of the general setting extends the findings of Rochet (1987) to interdependent value settings. Specifically, we identify a suitable generalization of the cyclical monotonicity condition from Rochet (1987), which serves as the exact characterization of ESP implementability, similar to how standard cyclical monotonicity characterizes strategy-proof implementation in the special case of private values. Additionally, we also derive a generalization of the one-dimensionality condition for interdependent value settings and demonstrate that, under this condition, generalized cyclical monotonicity is equivalent to monotonicity in the allocation function modulo types that are entangled. This finding implies that the analogous result in Myerson (1981) (modified for strategy-proof implementation) as a special case.

ESP implementation shares certain conceptual similarities with the problem of maxmin implementation studied in Tang and Zhang (2021). In both contexts, agents respond conservatively to the uncertainties they face within the mechanism. In our study, these uncertainties relate to the strategies and types of other agents, while in Tang and Zhang (2021), they arise from ambiguity in the execution of the mechanism. Both papers characterize implementation using a version of cyclical monotonicity, although the specific formulations and interpretations differ.

ESP implementation is also related to the literature on robust mechanism design, pioneered by Bergemann and Morris (2005). Brooks and Du (2024) examined the design of optimal mechanisms under both informational and strategic uncertainties, focusing on the worst-case performance of a mechanism across all possible information structures and equilibria. Similarly, our work considers implementation in ex post dominant strategy equilibrium, addressing both types of uncertainties simultaneously.

2 Model

A finite group $\{1, \ldots, n\}$ of agents must collectively choose from a set X of alternatives. Each agent i possesses a piece of private information $\theta_i \in \Theta_i$, known as her *type*, where the set of her types Θ_i is assumed to be finite. The set of all type profiles of the agents is denoted as $\Theta := \Theta_1 \times \ldots \times \Theta_n$.

Agent i's utility function is given by $u_i(x;\theta) - t_i$, where $u_i : X \times \Theta \to \mathbb{R}$ determines the value of each alternative x to agent i given each type profile θ, and t_i is the transfer amount she pays to the mechanism. Note that u_i can also depend on other agents' types θ_{-i}, making the environment one of interdependent values.

---

[4] Although we show there that *efficient* social choice functions may not be ISP implementable, which is also noted by Williams and Radner (1988).

As is standard, a mechanism (with transfers) consists of sets of messages $M_1, ..., M_n$, one for each agent, an outcome function $q : M_1 \times ... \times M_n \to X$, and a transfer function $\tau = (\tau_1, ..., \tau_n) : M_1 \times ... \times M_n \to \mathbb{R}^n$. Each agent i reports a message $m_i \in M_i$. The report profile $m = (m_1, ..., m_n)$ from all agents determine the chosen alternative $q(m)$ and the payment $\tau_i(m)$ from each agent i.

Given a mechanism, a strategy $\sigma_i \in \Delta(M_i)$ is said to be *ex post dominant* for agent i of type θ_i if her utility from playing σ_i is weakly higher than that from playing any alternative strategy $\sigma_i' \in \Delta(M_i)$ regardless of the messages reported by the other agents and the types of the other agents.

We focus on identifying which social choice functions—mappings from Θ to X—are consistent with some equilibrium of a mechanism in which every agent of every type has, and plays, an ex post dominant strategy. By the revelation principle[5], we can restrict attention to *direct mechanisms*, where $M_i = \Theta_i$ for every agent i.

Definition. A direct mechanism given by (q, τ) is *ex post strategy-proof (ESP)* if for any agent i of any type $\theta_i \in \Theta_i$, we have

$$u_i\Big(q(\theta_i, m_{-i}); \theta_i, \theta_{-i}\Big) - \tau_i(\theta_i, m_{-i}) \geq u_i\Big(q(m_i, m_{-i}); \theta_i, \theta_{-i}\Big) - \tau_i(m_i, m_{-i})$$

for any report $m_i \in \Theta_i$, any type profile $\theta_{-i} \in \Theta_{-i}$ and any report profile $m_{-i} \in \Theta_{-i}$ from the other agents.

Thus, a direct mechanism is ESP if truthtelling is an ex post dominant strategy for every agent. Note that if $u_i(x; \theta)$ does not vary with θ_{-i}, the environment is one of private values, and ESP reduces to standard strategy-proofness. In such cases, strategy-proofness is known to be equivalent to ex post incentive compatibility (EPIC). In contrast, when preference interdependence is present, ESP is generally stronger than EPIC, as EPIC only requires the inequality above to hold when other agents are truthful, i.e. when $m_{-i} = \theta_{-i}$.

A social choice function $q : \Theta \to X$ is *ESP-implementable*, or simply *implementable*, if there exists a transfer function τ such that the direct mechanism given by (q, τ) is ESP. Our objective is to characterize all social choice functions that are implementable.

3 Analysis

Our exposition will progress from specific to general, in order to build understanding incrementally. We begin with a simple auction example, then advance to a more general single-good auction model. Next, we explore a class of environments that are "one-dimensional", which includes auctions as a special case. Finally, we characterize implementation for the most general setting. All proofs are provided in the Appendix.

Leading example

We begin with a simple example to illustrate some key insights. The arguments in this section will be somewhat informal, but they will be formalized in subsequent sections.

---

[5] The proof of which follows standard argument and is therefore omitted.

A seller is looking to sell a single indivisible good, and there are two interested buyers, agents $i = 1, 2$. Each agent i has a type represented by a real number θ_i, which can take any value from $\Theta_i = \{0, \frac{1}{K}, \frac{2}{K}, \cdots, \frac{K}{K}\}$, where K is a positive natural number. Observe that the types are evenly distributed between 0 and 1, with K acting as a parameter that reflects the density of the types. As K increases, the gap between two neighboring types $(1/K)$ decreases, resulting in a more densely distributed set of types. As K approaches infinity, Θ_i approximates a continuum.

Agent i's valuation for the good is given by $v_i(\theta_i, \theta_j) = \theta_i + \beta\theta_j$, where $\beta \in [0, 1]$ is a parameter that reflects the degree of preference interdependence. A higher value of β indicates that the other agent's type has a greater influence on how one evaluates the good. Specifically, when $\beta = 0$, the model represents a pure private value auction, whereas when $\beta = 1$, the model represents a pure common value auction.

A direct mechanism is given by an allocation function $q = (q_1, q_2) : \Theta \to \Delta^2$ [6] and a transfer function $\tau = (\tau_1, \tau_2) : \Theta \to \mathbb{R}^2$. Here, $q_i(m)$ represents the probability that agent i wins the good based on the report profile m, and $\tau_i(m)$ denotes her payment. If the actual type profile is θ, then agent i's utility is $q_i(m)v_i(\theta) - \tau_i(m)$.

Which allocation functions are implementable? We will show that the answer depends on the interplay between the two parameters, K and β. Specifically, if $\beta K > 1$, then q is implementable if and only if each q_i is constant in θ_i. Conversely, if $\beta K \leq 1$, then q is implementable if and only if each q_i is (weakly) increasing in θ_i.

To see why, let us focus on an agent i and two of her neighboring types, θ_i and $\theta_i' = \theta_i - 1/K$. Incentive compatibility requires that agent i does not want to report type θ_i' when her type is θ_i, regardless of the actual type θ_j and the report m_j from the other agent. Formally, for $\forall \theta_j, m_j \in \Theta_j$:

$$q_i(\theta_i, m_j)v_i(\theta_i, \theta_j) - \tau_i(\theta_i, m_j) \geq q_i(\theta_i', m_j)v_i(\theta_i, \theta_j) - \tau_i(\theta_i', m_j). \tag{1}$$

Similarly, for type θ_i' to have no incentive to report to be type θ_i, we require $\forall \theta_j, m_j \in \Theta_j$:

$$q_i(\theta_i', m_j)v_i(\theta_i', \theta_j) - \tau_i(\theta_i', m_j) \geq q_i(\theta_i, m_j)v_i(\theta_i', \theta_j) - \tau_i(\theta_i, m_j). \tag{2}$$

The usual argument (by adding the two inequalities and rearranging) leads to $q_i(\theta_i, m_j) \geq q_i(\theta_i', m_j)$. This establishes a familiar monotonicity condition: the winning probability must be weakly increasing in one's type.

Can we make $q_i(\theta_i, m_j)$ *strictly* greater than $q_i(\theta_i', m_j)$? To answer this question, we first express the inequalities (1) and (2) across every $\theta_j \in \Theta_j$. From inequality (1), we have: $\forall m_j \in \Theta_j$,

$$\Big(q_i(\theta_i, m_j) - q_i(\theta_i', m_j)\Big) \min_{\theta_j \in \Theta_j} v_i(\theta_i, \theta_j) \geq \tau_i(\theta_i, m_j) - \tau_i(\theta_i', m_j). \tag{3}$$

Similarly, from inequality (2): $\forall m_j \in \Theta_j$,

$$\tau_i(\theta_i, m_j) - \tau_i(\theta_i', m_j) \geq \Big(q_i(\theta_i, m_j) - q_i(\theta_i', m_j)\Big) \max_{\theta_j \in \Theta_j} v_i(\theta_i', \theta_j). \tag{4}$$

---

[6] Δ^2 denotes the 2-simplex: the set of pairs $(p_1, p_2) \in [0, 1]^2$ such that $\sum_i p_i \leq 1$.

It is clear that $q_i(\theta_i, m_j)$ can be strictly greater than $q_i(\theta'_i, m_j)$ for some m_j only if [7]

$$\min_{\theta_j \in \Theta_j} v_i(\theta_i, \theta_j) \geq \max_{\theta_j \in \Theta_j} v_i(\theta'_i, \theta_j). \tag{5}$$

If inequality (5) is violated, then $q_i(\theta_i, m_j)$ must equal $q_i(\theta'_i, m_j)$ for any m_j. In other words, θ_i and θ'_i cannot lead to different winning probabilities for agent i in any ESP mechanism; thus, the two types are "*unscreenable*".[8] We conclude that inequality (5) is a necessary condition for the screenability of two types.

To interpret, although θ_i appears to be a "higher" type than θ'_i in the sense that, for any θ_j, the value of the good to agent i is greater under θ_i than under θ'_i, this alone is not sufficient to guarantee that type θ_i can have a strictly higher winning probability in some ESP mechanism. Instead, a more stringent condition is necessary: the lowest possible value of the good under θ_i must be no less than the highest possible value under θ'_i. Otherwise, θ_i and θ'_i must lead to the same winning probability.

Screenability will be shown to play a crucial role in ESP-implementation. To better understand when types are screenable and when they are not, let us consider the *valuation interval* for each type θ_i, defined as:

$$[\theta_i] := \left[\min_{\theta_j} v_i(\theta_i, \theta_j), \quad \max_{\theta_j} v_i(\theta_i, \theta_j)\right] (= [\theta_i, \theta_i + \beta] \text{ in the current example}).$$

The valuation interval represents the smallest closed interval in $\mathbb{R}$ that encompasses all possible ex post values that agent i may assign to the good given type θ_i. Since agent i knows her own type, she can ascertain that her value must fall within this interval, regardless of her beliefs about θ_j. It is important to note that the interval $[\theta_i]$ is not necessarily the set of all possible values that agent i may assign to the good given her type θ_i, as some values within the interval may not obtain for any θ_j.

Now, returning to inequality (5), which is necessary for screenability, we can visualize the inequality and its violation through plots:

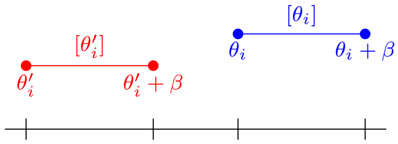

Inequality (5) satisfied

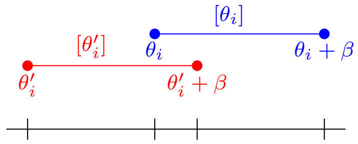

Inequality (5) violated

We observe that inequality (5) is violated (and consequently, the two types are unscreenable) if $[\theta_i]$ and $[\theta'_i]$ have a non-degenerate, i.e., nonempty and non-singleton, intersection. In this case, we say that θ_i and θ'_i *overlap*. Numerically, this occurs when $\beta K > 1$. Graphically, this occurs when the two valuation intervals widen to the extent that they fill the gap between them and begin to encroach upon one another. This indicates that

---

[7] This inequality resembles the largest and smallest payment rules in Kos and Messner (2013).

[8] We note that even though the two types are unscreenable from agent i's perspective, they may still lead to different winning probabilities for agent j. Such "non-self responsive" mechanisms may still be of interest in general. Niemeyer and Preusser (2024) shows that when agents are "informatinally small", for certain environments without transfers, a class of non-self-responsive mechanisms, termed as "ranking-based" mechanisms, achieve approximately optimality.

when there is excessive preference interdependence (reflected in a greater β) relative to the density of types (measured by K), neighboring types start to overlap, rendering them unscreenable.

Inequality (5) is satisfied when $[\theta_i]$ lies entirely to the right of $[\theta_i']$ without any non-degenerate intersection. In this case, θ_i might be screened from θ_i', that is, there could exist an implementable allocation function q such that $q_i(\theta_i, m_j) > q_i(\theta_i', m_j)$ for some m_j. However, unscreenablity is still a possibility. To illustrate this, suppose $\beta K > 1$, which implies that neighboring types overlap. Consider a pair of non-neighboring types θ_i and θ_i' such that $\theta_i > \theta_i'$, and assume that they do not overlap. Even so, since $\beta K > 1$, θ_i overlaps with $\theta_i - 1/K$, which in turn overlaps with $\theta_i - 2/K$, and so forth, until we eventually reach θ_i'. In other words, we can "travel" from θ_i to θ_i' through a finite sequence of types, each overlapping with, and therefore unscreenable from, the next. In this case, we say that the two types are *entangled*. Since unscreenability is transitive, two entangled types are unscreenable even if they do not overlap directly. This possibility is illustrated below, where we travel from θ_i to θ_i' via an intermediate type $\hat{\theta}_i$, which overlaps with both θ_i and θ_i':

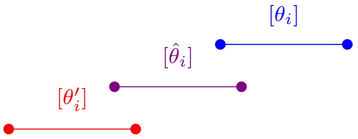

If, on the other hand, two types are not entangled, then they are screenable. In the current example, this occurs when $\beta K \leq 1$, as even neighboring types do not overlap. In this scenario, any allocation function q where q_i is non-decreasing in θ_i for each i can be implemented. The construction of the supporting transfer function follows a standard approach: for any m_j, we set $\tau_i(0, m_j) = 0$; for each $\theta_i = \frac{1}{K}, \ldots, \frac{K}{K}$, we recursively define:

$$\tau_i(\theta_i, m_j) = \tau_i(\theta_i - 1/K, m_j) + \theta_i \Big(q_i(\theta_i, m_j) - q_i(\theta_i - 1/K, m_j) \Big).$$

We summarize the lessons from this exercise with the following two observations:

- Monotonicity: If two types of the same agent are not entangled, then they are screenable, and the higher type can lead to a higher winning probability for the agent.
- Screenability: If two types of the same agent are entangled, including the special case that they overlap, then they are not screenable.

Single-good auction

We turn our attention to single-good auctions in general, and demonstrate how the insights from the leading example can be extended. Now there are n buyers, agents $i = 1, ..., n$, who are interested in a single indivisible good. Each agent i's type $\theta_i \in \Theta_i$ need not be a real number. In particular, a type can be "multi-dimensional", i.e. a vector of real numbers.

Agent i's valuation for the good is given by a function $v_i : \Theta \to \mathbb{R}$. We assume that each agent's valuation always vary with the types of other agents. Formally, what this entails is that for any agent i of any type $\theta_i \in \Theta_i$, there exist some $\theta_{-i}, \theta'_{-i} \in \Theta_{-i}$ such that $v_i(\theta_i, \theta_{-i}) \neq v_i(\theta_i, \theta'_{-i})$. [9]

A direct mechanism is given by an allocation function $q = (q_1, ..., q_n) : \Theta \to \Delta^n$ [10] and transfer function $\tau = (\tau_1, ..., \tau_n) : \Theta \to \mathbb{R}^N$.

We will show that the two observations from leading example continue to hold in this more general setting. In fact, we will show that the two conditions, monotonicity and screenability, jointly characterize implementability.

In the current setting, the valuation interval of a type $\theta_i \in \Theta_i$ is

$$[\theta_i] := \left[\min_{\theta_{-i}} v_i(\theta_i, \theta_{-i}), \quad \max_{\theta_{-i}} v_i(\theta_i, \theta_{-i})\right].$$

Same as in the leading example, two types of the same agent are said to overlap if their respective valuation intervals have a non-degenerate intersection, whereas two types are said to be entangled if they can be linked by a finite sequence of types, one overlapping the next.

In addition, we say that θ_i is *higher than* θ'_i if (1) the two types are not entangled, and (2) $[\theta_i]$ lies to the right of $[\theta'_i]$ on the real line, i.e. $\max[\theta'_i] \leq \min[\theta_i]$. The "monotonicity" observation in the leading example is the same as saying that if a type is higher than another, the higher type can lead to a higher winning probability in some ESP mechanism.

The observations from the leading example can be generalized and formalized as follows:

Proposition 3.1. *An allocation function q is implementable if and only if for any agent i, any pair of her types $\theta_i, \theta'_i \in \Theta_i$ and any report profile $m_{-i} \in \Theta_{-i}$ from the other agents, we have:*

1. *Monotonicity: If θ_i is higher than θ'_i, then $q_i(\theta_i, m_{-i}) \geq q_i(\theta'_i, m_{-i})$;*
2. *Screenability: If θ_i and θ'_i are entangled, then $q_i(\theta_i, m_{-i}) = q_i(\theta'_i, m_{-i})$.*

The proposition is a corollary of a more general result that will be presented in the next section. The heuristic analysis provided in the leading example offers a rough outline of the proof.

One important implication of Proposition 3.1 is that, as illustrated in the leading example, the scope of ESP implementation is jointly constrained by the density of types, i.e., how closely the valuation intervals are situated to one another, and the extent of preference interdependence, i.e., the widths of those valuation intervals. When types are densely distributed or when preference interdependence is significant, the valuation intervals will begin to overlap, rendering types unscreenable and limiting the ability of an allocation function to differentiate between them. If Θ_i is a continuum and v_i is

---

[9] This assumption primarily serves to simplify exposition. In the next section, we will discuss a more general setting that does not impose such a restriction.

[10] Δ^n denotes the n-simplex: the set of $(p_1, \cdots, p_n) \in [0,1]^n$ such that $\sum_i p_i \leq 1$.

continuous in θ_i, it becomes evident that as long as there is "minimal" preference interdependence in the sense that the the widths of valuation intervals do not vanish, then no two types can be screened. [11] Consequently, agent i's winning probability cannot depend on her own report.

Conversely, when types are not too densely distributed, a certain degree of preference interdependence would not hinder implementation as long as preference interdependence does not become so great as to cause valuation intervals to overlap with each other. This is exemplified in the leading example when $\beta K \leq 1$, where every allocation function that ensures a non-decreasing winning probability based on one's own report is implementable. Jehiel, Meyer-ter Vehn, Moldovanu, and Zame (2006) study auctions where types are drawn from a continuum and show a generic impossibility for the existence of nontrivial ex post incentive compatible (EPIC) mechanisms when types are multi-dimensional. In light of our findings, this impossibility critically hinges on the assumption of continuous types. In contrast, if types are not as densely distributed, then nontrivial ESP mechanisms—which are *a fortiori* also EPIC—can exist even when types are multi-dimensional.

Next, we investigate whether it is possible to implement efficient allocation functions, i.e., those that allocate the good to an agent who values it the highest among all agents. For simplicity, we assume that no two agents assign the same value to the good; that is, for any $\theta \in \Theta$, $v_i(\theta) \neq v_j(\theta)$ if $i \neq j$. Let $i^*(\theta)$ denote the unique agent who assigns the highest value to the good. Consequently, there is a unique efficient allocation function q^E, which is defined as follows:

$$q_i^E(\theta) = \begin{cases} 1 & \text{if } i = i^*(\theta) \\ 0 & \text{otherwise.} \end{cases}$$

It is well known that q^E can be implemented by an EPIC mechanism if the underlying environment satisfies a condition known as "single-crossing" (Maskin (1992), Dasgupta and Maskin (2000)). Essentially, the single-crossing condition stipulates that if an agent assigns the highest value to the good among all agents, then she will continue to do so if her type (which is assumed to be a real number in the related works) increases while the types of the other agents remain unchanged.

We will show that a stronger, yet conceptually similar, condition is necessary and sufficient to ESP-implement q^E. For two types θ_i and θ_i' of agent i, we say that θ_i' is *weakly higher* than θ_i if the two types are entangled or if θ_i' is higher than θ_i.

Definition. A single-good auction environment is said to satisfy *strong single-crossing* if for any agent i, whenever $i = i^*(\theta_i, \theta_{-i})$ for some $\theta_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$, then $i = i^*(\theta_i', \theta_{-i})$ for any $\theta_i' \in \Theta_i$ that is weakly higher than θ_i

To draw a direct comparison between single-crossing and strong single-crossing, we assume, as Dasgupta and Maskin (2000) does, that types are real numbers and that v_i is increasing in θ_i. Single-crossing requires that if $i = i^*(\theta)$, then $i = i^*(\theta_i', \theta_{-i})$ for any $\theta_i' > \theta_i$. Strong single-crossing requires that $i = i^*(\theta_i', \theta_{-i})$ not only for any θ_i'

---

[11] Since we technically maintain the assumption that Θ_i is finite for every agent i, a more rigorous statement would involve a sequence of environments, each with a finite type space, converging to one with a continuum of types, as demonstrated in the leading example where $K \to \infty$ while β does not vanish.

greater than θ_i, but also for θ_i' smaller than θ_i as along as the two types are entangled. Essentially, what strong single-crossing entails is that for every agent i, the types of that agent which make her the one who values the good the highest under some θ_{-i} must be screenable from the types that do not. In our leading example, when $\beta K > 1$, all types of the same agent are entangled, and hence strong single-crossing necessitates that whether an agent is the one who values the good the highest does not depend on her type, which is clearly violated. In fact, it is easy to verify that in that example, strong single-crossing is satisfied if $\beta \leq 1/K$, while single-crossing is satisfied if $\beta \leq 1$.

Proposition 3.2. *The efficient allocation function q^E is implementable if and only if the environment satisfies strong single-crossing.*

It is easy to see why strong single-crossing is necessary: By Proposition 3.1, once agent i wins the good under reported type profile θ, she must continue to win the good if her type becomes weakly higher. Since under q^E an agent wins the good whenever she is the one who values it the highest, for q^E to be implementable, it is required that when an agent is the one who values the good the highest, she must continue to be the highest valuer when her type becomes weakly higher, which is exactly what strong single-crossing entails.

Sufficiency is established by constructing supporting transfers as follows:

- If agent i does not win, she does not pay anything.
- If agent i wins as she reports θ_i and others report m_{-i}, she pays the value of the following maximization problem:

$$\max_{\tilde{\theta} \in \Theta} v_i(\tilde{\theta}) \qquad \text{s.t.} \quad q_i^E(\tilde{\theta}_i, m_{-i}) = 0.$$

Observe that losing agents do not pay, while the winning agent pays an amount equal to the highest possible value she would assign to the good, assuming she had a type that would not lead her to win, where this hypothetical value is determined entirely by other agents' reports. The rationale behind this transfer scheme is analogous to the one presented in Dasgupta and Maskin (2000), which is used to ex post implement q^E, and is conceptually similar to the Vickrey auction.

In contrast, Williams and Radner (1988) shows that when each agent's types form an interval on the real line, implementing q^E by an ISP mechanism is impossible, even with minimal preference interdependence. This impossibility, in light of our findings, again hinges on the assumption of continuous types. As shown in Proposition 3.2 and the leading example where $\beta K \leq 1$, if types are less densely distributed, the presence of mild preference interdependence does not necessarily preclude the implementability of q^E by an ESP, and henceforth ISP, mechanism.

Generalized one-dimensional environments

Single-good auctions are a special case of a broader class of environments in which monotonicity and screenability jointly characterize implementability. We refer to this class as *generalized one-dimensional* because it includes standard one-dimensional environments when there is no preference interdependence.[12]

We first define this class of environments. For each agent i, we use

$$u_i^{xy}(\theta) := u_i(x;\theta) - u_i(y;\theta)$$

to denote the marginal value of alternative x over alternative y given type profile θ, and use

$$[\theta_i]^{xy} := \left[\min_{\theta_{-i}} u_i^{xy}(\theta_i,\theta_{-i}), \quad \max_{\theta_{-i}} u_i^{xy}(\theta_i,\theta_{-i})\right]$$

to denote the closed interval bounded by the minimum and maximum marginal values of x over y conditional on type θ_i. If agent i knows her own type, she can determine that the marginal value of x over y must lie within this interval, regardless of her beliefs about θ_{-i}. We refer to such $[\theta_i]^{xy}$ as a *marginal value interval*.

The valuation intervals we introduced for single-good auctions differ from marginal value intervals, although they are closely related. A valuation interval depends only on the type, while a marginal value interval depends on both the type and the pair of alternatives involved. In single-good auctions, an alternative in set X can be represented as a vector of winning probabilities $p = (p_1, ..., p_n)$. The payoff to agent i from any $p \in X$ is given by $u_i(p;\theta) = p_i v_i(\theta)$. Therefore, for any $p, p' \in X$, we have $u_i^{pp'}(\theta) = (p_i - p_i')v_i(\theta)$. The marginal value interval of type θ_i regarding p vs. p' is thus

$$[\theta_i]^{pp'} = \begin{cases} \left[(p_i - p_i')\min_{\theta_{-i}} v_i(\theta_i,\theta_{-i}), \quad (p_i - p_i')\max_{\theta_{-i}} v_i(\theta_i,\theta_{-i})\right] & \text{if } p_i \geq p_i' \\ \\ \left[(p_i - p_i')\max_{\theta_{-i}} v_i(\theta_i,\theta_{-i}), \quad (p_i - p_i')\min_{\theta_{-i}} v_i(\theta_i,\theta_{-i})\right] & \text{if } p_i < p_i'. \end{cases}$$

In other words, $[\theta_i]^{pp'}$ can be obtained from $[\theta_i] = [\min_{\theta_{-i}} v_i(\theta_i,\theta_{-i}), \max_{\theta_{-i}} v_i(\theta_i,\theta_{-i})]$ by multiplying each element of the latter by the constant $p_i - p_i'$. A useful observation is that θ_i and θ_i' overlap if and only if $[\theta_i]^{pp'}$ and $[\theta_i']^{pp'}$ have a non-degenerate intersection for any $p, p' \in X$ where $p_i \neq p_i'$.

The upcoming analysis will rely on examining the relative "positions" of different types' marginal value intervals on the real line. Given two closed intervals A and B, we define:

- $A \to B$ if $a \leq b$ for every $a \in A$ and $b \in B$, with $a < b$ for some $a \in A$ and $b \in B$;
- $A \bowtie B$ if $A \not\to B$ and $B \not\to A$.

The notation $A \to B$ indicates that interval B is positioned entirely to the right of interval A and they do not have a non-degenerate intersection. Specifically, one of the following scenarios must hold, depending on whether one or both intervals are degenerate, and whether they intersect at the endpoints.

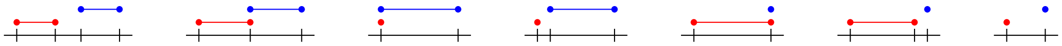

$A \to B$

---

[12] For a definition of standard one-dimensional environments, see Section 5.6 of Börgers (2015).

If $A \bowtie B$, then there are three possibilities:

1. If both A and B are non-degenerate, they have a non-degenerate intersection.

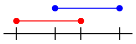

2. If one of A or B is degenerate and the other is non-degenerate, then the degenerate interval lies strictly within the non-degenerate one.

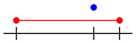

3. If both A and B are degenerate, they coincide.

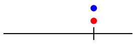

Now we formally define generalized one-dimensional environments.

Definition. *(Generalized one-dimensionality)* An environment is said to satisfy *generalized one-dimensionality (G1D)* if, for each agent i, there exist complete and transitive relations $\succsim_i^X$ on X (with $\succ_i^X$ and $\sim_i^X$ representing the induced strict and equivalence relations respectively) and $\succsim_i^\Theta$ on Θ_i (with $\succ_i^\Theta$ and $\sim_i^\Theta$ representing the induced strict and equivalence relations respectively), such that:

1. If $x \succ_i^X y$ and $\theta_i \succ_i^\Theta \theta_i'$, then $[\theta_i']^{xy} \to [\theta_i]^{xy}$.
2. If $x \sim_i^X y$, then $[\theta_i]^{xy} = \{0\}$.
3. If $\theta_i \sim_i^\Theta \theta_i'$ and $\theta_i \neq \theta_i'$, then both of the following conditions hold:
   (a) There exists a finite sequence of agent i's types $(\theta_i^1, \dots, \theta_i^k)$ such that $\theta_i^1 = \theta_i$, $\theta_i^k = \theta_i'$, and for any $x \not\sim_i^X y$, we have $[\theta_i^\kappa]^{xy} \bowtie [\theta_i^{\kappa+1}]^{xy}$ for $\kappa = 1, \dots, k-1$.
   (b) $[\theta_i]^{xy} \cup [\theta_i']^{xy}$ is not a singleton for any $x \not\sim_i^X y$.

If these conditions are met, we say that the environment *manifests generalized one-dimensionality in* $(\succsim_i^X, \succsim_i^\Theta)_{i=1:n}$.

To better understand this definition, it is useful to recall what makes an environment one-dimensional in private value settings. There, one-dimensionality means that for each agent, alternatives and her types can be ranked from "high" to "low", under which the agent's utility function exhibits increasing differences: If one alternative is ranked higher than another, the marginal value of the higher-ranked alternative over the lower-ranked one is greater to the agent if she has a higher-ranked type than if she has a lower-ranked type. If the two alternatives are equally ranked, they are of equal value to the agent regardless of her type.

Now, with interdependent values, G1D also requires that for each agent i, both the alternatives and her types can be ranked from high to low according to $\succsim_i^X$ (the ranking of alternatives) and $\succsim_i^\Theta$ (the ranking of types). These rankings ensure that the agent's utility function exhibits a property similar to increasing differences. Specifically, Condition (1) states that the marginal value of a higher-ranked alternative x over a lower-ranked alternative y is greater for agent i if she has a higher-ranked type θ_i compared to when she has a lower-ranked type θ_i', in the sense that the minimal possible marginal value of x over y given θ_i must be at least as large as the maximal possible marginal value of x over y given θ_i'. Condition (2) requires that if two alternatives are equally ranked, they are of equal value for the agent, regardless of her type.

Condition (3) is specific to interdependent value settings and addresses the scenario where a pair of types may be equally ranked due to "entanglement". This phenomenon has been observed in our leading example and in single-good auctions. According to Condition (3a), entanglement occurs when two types can be linked through a finite sequence of types. In this sequence, the marginal value intervals of consecutive types regarding any two non-equally ranked alternatives must have a non-degenerate intersection. In single-good auctions, this is indeed the case for two types that are entangled. For the sequence of overlapping types that link them, the marginal value intervals of any neighboring types along this sequence regarding any p vs. p' where $p_i \neq p_i'$ (which is shown to imply that p and p' are not equally ranked in the proof of Proposition 3.1) have a non-degenerate intersection. We will show that, similar to the leading example and in single-good auctions, entanglement will imply unscreenability.

Condition (3b) is a mild restriction that primarily aims to simplify the exposition. To understand what this condition rules out, suppose $[\theta_i]^{xy} \cup [\theta_i']^{xy}$ is a singleton for some x and y that not are equally ranked, i.e. $x \not\sim_i^X y$. In this case, both $[\theta_i]^{xy}$ and $[\theta_i']^{xy}$ must also be singletons, indicating that preference interdependence is absent at either type θ_i or θ_i' when comparing x and y, since the marginal value of x over y would be the same regardless of the types of other agents. Furthermore, θ_i and θ_i' are "redundant" with respect to the comparison of x and y, as they assign the same marginal value to x over y. Thus, Condition (3b) effectively rules out the redundancy of "private value" types. [13]

In private value settings, types cannot be entangled, because the marginal value interval of any type regarding any pair of alternatives is degenerate. Thus, once redundant types are assumed away (to avoid the slight complication due to a violation of Condition (3b)), standard one-dimensionality is equivalent to G1D.

Single-good auctions satisfy G1D. Specifically, a single-good auction environment manifests G1D[14] in the following rankings of alternatives

$$p \succsim_i^X p' \quad \text{if } p_i \geq p_i';$$

and types

$$\theta_i \succsim_i^\Theta \theta_i' \quad \text{if } \theta_i \text{ is weakly higher than } \theta_i'.$$

It follows $\theta_i \sim_i^\Theta \theta_i'$ if θ_i and θ_i' are entangled, or $\theta_i \succ_i^\Theta \theta_i'$ if θ_i is higher than θ_i'. Therefore, Proposition 3.1 can be reinterpreted to state that in single-good auctions, an allocation function is implementable if and only if a higher-ranked type leads to a higher or equally-ranked alternative, while equally ranked types lead to equally ranked alternatives. In fact, this principle holds true more generally for any environment that satisfies

---

[13] If G1D is defined without Condition (3b), the subsequent analysis would still hold, but the statement of Proposition 3.3 would become more complex. This complexity arises because redundant "private value" types would imply multiple rankings of types, all of which are compatible with what is stipulated by G1D.

[14] This is formally established in the proof of Proposition 3.1 in the Appendix.

G1D:

Proposition 3.3. *Suppose the environment manifests generalized one-dimensionality in* $(\succsim_i^X, \succsim_i^\Theta)_{i=1:n}$. *A social choice function* q *is implementable if and only if for any agent* i, *any two of her types* $\theta_i, \theta_i' \in \Theta_i$ *and any report profile* $m_{-i} \in \Theta_{-i}$ *from the other agents, we have:*

1. *Monotonicity: If* $\theta_i \succ_i^\Theta \theta_i'$, *then* $q_i(\theta_i, m_{-i}) \succsim_i^X q_i(\theta_i', m_{-i})$;
2. *Screenability: If* $\theta_i \sim_i^\Theta \theta_i'$, *then* $q_i(\theta_i, m_{-i}) \sim_i^X q_i(\theta_i', m_{-i})$.

Implementability in general

In this section, we will examine the most general environment as described in the Model section. Specifically, we will present two characterizations for social choice functions that are implementable.

First characterization: generalized cyclical monotonicity

The first characterization extends the well-known cyclical monotonicity condition[15] to interdependent value settings. This condition is defined as follows:

Definition. A social choice function q is said to satisfy *generalized cyclical monotonicity (GCM)* if for every agent i, every report profile $m_{-i} \in \Theta_{-i}$ from the other agents, and every finite sequence of agent i's types $(\theta_i^1, ..., \theta_i^k)$ such that $\theta_i^1 = \theta_i^k$, we have

$$\sum_{\kappa=1}^{k-1} \max_{\theta_{-i}} \left\{ u_i\Big(q(\theta_i^{\kappa+1}, m_{-i}); \theta_i^\kappa, \theta_{-i}\Big) - u_i\Big(q(\theta_i^\kappa, m_{-i}); \theta_i^\kappa, \theta_{-i}\Big) \right\} \leq 0.$$

To interpret this, the term in the curly brackets reflects the maximum utility gain (modulo transfers) for agent i of type θ_i^κ when, instead of reporting her true type, she deviates to report $\theta_i^{\kappa+1}$. GCM is satisfied if, for any closed chain of such deviations, these "best-case scenario" utility gains do not sum to a positive number, regardless of how other agents report.

Note that if u_i does not vary with θ_{-i}, as is the case with private values, the $\max$ operator can be omitted, making GCM equivalent to the standard cyclical monotonicity condition. Just as standard cyclical monotonicity characterizes strategy-proof implementability in private value settings (Rochet (1987)), GCM characterizes ESP implementability in the more general context with interdependent values:

Proposition 3.4. *A social choice function* q *is implementable if and only if it satisfies generalized cyclical monotonicity.*

The reason standard cyclical monotonicity generalizes in this way to characterize ESP implementability is rooted in the requirement for ex post optimality in the presence of interdependent preferences. When contemplating a deviation, an agent evaluates all possible type profiles of others. Incentive compatibility must therefore ensure that such a deviation is not profitable, even under the most favorable type profile of others—hence

---

[15] Standard cyclical monotonicity originates from Rockafellar (1970) and has been applied to mechanism design since Rochet (1987).

the inclusion of the $\max$ operator. It is important to note that this $\max$ operator is included within the summation across the deviation chain. This reflects the fact that a deviation chain is broken down into a sequence of deviations, each considered separately and each potentially involving different "most favorable" θ_{-i} profiles.

Second characterization: optimal matching

The second characterization is quite different in nature. It draws an interesting connection between social choice functions and matching problems, showing that implementability can be characterized in terms of optimal matchings. Such a connection has been noted in Rahman (2024) and Dworczak and Zhang (2017). Specifically, they demonstrate that in a single-agent screening problem, implementability of a social choice function is equivalent to the utilitarian efficiency of an induced matching between the different types of that agent and the alternatives, treating each type as if it were an individual. Our characterization will extend their findings to settings with multiple agents and interdependent values.

To set the stage, fix a social choice function q, an agent i, and a report profile m_{-i} from the other agents. We can view the function $q(\cdot, m_{-i})$ as defining a matching $\mu : \Theta_i \to X$ between types of agent i and the alternatives in X. Specifically, each type θ_i of agent i is matched to an alternative given by $\mu(\theta_i) := q(\theta_i, m_{-i})$. In this framework, we treat the different types of the same agent as distinct agents, with μ describing a matching between these agents and the alternatives. We call such a matching μ as a *type-alternative matching for agent i*.

Given any type-alternative matching $\mu : \Theta_i \to X$ for agent i, another type-alternative matching for agent i, $\mu' : \Theta_i \to X$, is called a *reshuffling* of μ if there exists a permutation (self-bijection) $\pi : \Theta_i \to \Theta_i$ such that $\mu'(\theta_i) = \mu(\pi(\theta_i))$. Note that μ' is obtained from μ by reassigning to each θ_i with the alternative originally assigned to $\pi(\theta_i)$.

Suppose that μ is a type-alternative matching for agent i, and μ' is a reshuffling of μ. Define

$$D_i(\mu', \mu) := \sum_{\theta_i \in \Theta_i} \max_{\theta_{-i}} \Big\{ u_i\Big(\mu'(\theta_i); \theta_i, \theta_{-i}\Big) - u_i\Big(\mu(\theta_i); \theta_i, \theta_{-i}\Big) \Big\}.$$

To understand the meaning of $D_i(\mu', \mu)$, note that the term in the curly brackets represents, for a given type profile θ_{-i} of the other agents, the net change in utility for type θ_i when the matching shifts from μ to μ'. The $\max$ operator indicates that the most favorable scenario for μ' is considered, while the $\sum$ operator aggregates these optimistically evaluated utility changes across different types. Therefore, $D_i(\mu', \mu)$ captures the most favorable assessment of efficiency improvement when comparing μ' against μ.

If $D_i(\mu', \mu)$ is not positive, we can conclude that, from a utilitarian perspective, there is no gain from reshuffling the matching from μ to μ', regardless of the subjective beliefs each θ_i holds about θ_{-i}. If it is impossible to reshuffle a matching to improve efficiency in this sense, we define the matching as reshuffling-efficient:

Definition. A type-alternative matching μ for agent i is *reshuffling-efficient* if $D_i(\mu', \mu) \leq 0$ for any reshuffling μ' of μ.

Furthermore, a social choice function q is said to be *reshuffling-efficient* if for any agent

i and report profile m_{-i} from the other agents, the induced type-alternative matching for agent i, $q(\cdot, m_{-i})$, is reshuffling-efficient.

Reshuffling-efficiency embodies a stringent notion of utilitarian optimality for the matching. We establish that this optimality is both necessary and sufficient for the corresponding social choice function to be implementable:

Proposition 3.5. *A social choice function q is implementable if and only if it is reshuffling-efficient.*

This result arises from the observation that any reshuffling of a type-alternative matching for agent i corresponds to a chain of deviations among her types, and vice versa. Consequently, the absence of improvement from reshuffling is equivalent to the condition that no chain of deviations can lead to an overall utility gain, which is precisely what GCM entails.

Note that if u_i does not vary with θ_{-i}, as is the case in private value settings, the $\max$ operator can be omitted. In this context, $D_i(\mu', \mu) \leq 0$ indicates that the utilitarian welfare under μ is weakly higher than the utilitarian welfare under any reshuffling of μ. Therefore, reshuffling-efficiency aligns with utilitarian efficiency as described in Rahman (2024) and Dworczak and Zhang (2017) within a single-agent setting, which is inherently one of private values. Consequently, Proposition 3.5 can be seen as a generalization of their findings to multi-agent, interdependent value settings.

An implication of Rahman (2024) and Dworczak and Zhang (2017) is that, in private value settings, any social choice function can be transformed into an implementable social choice function by simply reshuffling the type-alternative matching for each agent.[16] This is because, given any type-alternative matching μ for any agent i, it is always possible to find a reshuffling μ^* that maximizes the utilitarian welfare among reshufflings of μ. Furthermore, this implies that a nontrivial implementable social choice function always exists. However, both conclusions do not necessarily hold in interdependent value settings. To see that, consider the following example: For each agent i, each type θ_i, and any ordered pair of distinct alternatives x and y, suppose there exists a type profile $\theta_{-i}^{\theta_i, x, y} \in \Theta_{-i}$ such that

$$u_i(x; \theta_i, \theta_{-i}^{\theta_i, x, y}) - u_i(y; \theta_i, \theta_{-i}^{\theta_i, x, y}) > 0.$$

It is straightforward to see that for any non-constant type-alternative matching μ for agent i, and any reshuffling μ' of μ, we have $D_i(\mu', \mu) > 0$. As a result, no non-constant matching is reshuffling-efficient. Consequently, no implementable social choice function can be derived from a nontrivial social choice function through reshuffling. This implies that no nontrivial implementable social choice function exists in this scenario.

---

[16] See Corollary 2 of Dworczak and Zhang (2017).

4 Discussion

Iterated elimination of dominated strategies

In this paper, we focused on implementation in ex post dominant strategies. Naturally, one might wonder if relaxing this criterion to iterated elimination of dominated strategies could broaden the scope of implementation. Intuitively, such a relaxation would allow agents to narrow down possible type-strategy correlations of other agents, thereby reducing incentive constraints. The following example illustrates this possibility.

Consider a single-good auction with two agents. Imagine that the good is a piece of artwork. Agent 1 is an art dealer who knows and only cares about the investment value of the artwork, denoted as $\theta_1 \in \{0, 0.5, 1\}$. Agent 2 is a collector who cares about the investment value θ_1 of the artwork, and also derives a private enjoyment value, $\theta_2 \in \{0, 0.5, 1\}$, if she owns it. Specifically, we assume that:

$$v_1(\theta_1, \theta_2) = \theta_1, \qquad v_2(\theta_1, \theta_2) = \begin{cases} \theta_2 & \text{if } \theta_1 = 0 \text{ or } 0.5 \\ \theta_1 + \theta_2 & \text{if } \theta_1 = 1. \end{cases}$$

To interpret agent 2's valuation for the artwork, imagine that she receives an instantaneous utility of θ_2 from owning it. Furthermore, if there is little investment value in the artwork ($\theta_1 = 0$ or 0.5), she would prefer to keep it permanently. However, if there is a significant investment value ($\theta_1 = 1$), she will resell it in the future for θ_1.[17] Suppose the mechanism designer is the seller who aims to maximize expected sales revenue.

Observe that $[\theta_2] = [\theta_2, \theta_2 + 1]$. It is easy to verify that $[\theta_2 = 0] \bowtie [\theta_2 = 0.5] \bowtie [\theta_2 = 1]$. Therefore all types of agent 2 are entangled, and it follows from Proposition 3.1 that agent 2's report cannot unilaterally influence her own winning probability in an ESP mechanism. In this case, the best the seller can do is to focus solely on agent 1, turning the problem into a single-agent screening situation. It is well known that the revenue-maximizing mechanism is a posted-price mechanism where the optimal price p^* depends on the auctioneer's belief about the distribution of θ_1. Assume, for example, $p^* = 0.5$.

Now we show that if the seller believes the agents can perform iterated elimination of dominated strategies, the following mechanism can yield a higher revenue:

- If agent 1's reported type is 0.5 or 1, agent 1 wins the good and pays 0.5, while agent 2 pays nothing.
- If agent 1's reported type is 0, agent 2 has a chance to win the good: she wins if her reported type is 0.5 or 1. In this case, agent 1 pays nothing, and agent 2 pays 0.5.
- If both agents' reported types are 0, no one wins or pays anything.

Clearly, for agent 1, reporting to be type 0 when her true type is 1 is strictly dominated. After eliminating this dominated strategy for agent 1, agent 2 realizes that when her

---

[17] This interpretation ignores discounting and assumes an instant rather than continuation utility from owning the artwork. The utility function could be modified for realism, but that would beat our purpose of offering a simple illustration.

own report matters (i.e., when agent 1's report is 0), her valuation does not depend on agent 1's type, since agent 1's true type must be either 0 or 0.5. In this scenario, it becomes agent 2's dominant strategy to report truthfully when facing the fixed price of 0.5. Compared to the optimal ESP mechanism, which allows the seller to sell only to agent 1 at price 0.5, this modified mechanism provides the seller with the opportunity to sell to agent 2 at 0.5 if agent 1 does not make a purchase. The mechanism can be implemented as a sequential selling procedure, where the good is first offered to agent 1 at price 0.5, and, upon rejection, is then offered to agent 2 at the same price.

No-transfer

What if transfers cannot be used? Technically, no-transfer is equivalent to implementation with supporting transfer function τ such that $\tau_i(\theta) = 0$ for any θ. Implementation becomes significantly more challenging when transfers are not available to alleviate incentive constraints. In this section, we explore the extent of this challenge. In particular, we focus on identifying the conditions under which implementation of nontrivial, i.e. non-constant, social choice functions is possible and evaluating how restrictive these conditions may be.

Agent i of type θ_i is said to *unambiguously prefer* x to y if $u_i(x; \theta_i, \theta_{-i}) \geq u_i(y; \theta_i, \theta_{-i})$ for any $\theta_{-i} \in \Theta_{-i}$. If agent i of type θ_i unambiguously prefers x to y, her preference regarding this pair of alternatives is *ordinally* independent of the types of other agents—she knows with certainty that x is weakly better than y. The following lemma demonstrates that incentive constraints can be formulated in terms of unambiguous preferences.

Lemma 4.1. *In the no-transfer case, a social choice function q is implementable if and only if any agent i of any type $\theta_i \in \Theta_i$ unambiguously prefers $q(\theta_i, m_{-i})$ to $q(\theta_i', m_{-i})$ for any $\theta_i' \in \Theta_i$ and any $m_{-i} \in \Theta_{-i}$.*

The lemma follows directly from the ESP incentive constraints by setting $\tau_i(\theta) = 0$ for all θ, so the proof is omitted. This lemma implies that truthful reporting must lead to agent i's unambiguous favorite alternative among the ones she can unilaterally induce, fixing the reports m_{-i} from the other agents. Therefore, the set of alternatives that she can unilaterally induce, given any report profile m_{-i} from the other agents, must satisfy the property that each of her types has an unambiguous favorite in that set. More precisely:

Definition. A set of alternatives $Y \subset X$ is said to be *decidable for agent i* if there exists a surjective *decision function* $c_i : \Theta_i \to Y$ such for any $\theta_i \in \Theta_i$, if agent i's type is $\theta_i \in \Theta_i$, she unambiguously prefers $c_i(\theta_i)$ to any $x \in Y$.

Clearly, every singleton set is decidable for every agent. However, for an agent to have a nontrivial impact on the social choice outcome, she must be able to unilaterally induce multiple alternatives. Since a social choice function is nontrivial if and only if at least one agent has a nontrivial impact on it, nontrivial implementation ultimately hinges on the existence of a non-singleton decidable set for some agent, as is stated in the following result:

Proposition 4.1. *There exists a nontrivial implementable social choice function if and only if some non-singleton $Y \subset X$ is decidable for some agent i.*

The core idea behind Proposition 4.1 is that a dictatorial social choice function—where only one agent determines the outcome—is the simplest to implement. This is because it eliminates the need for interactions between different agents' strategies, making it inherently strategy-proof. Consequently, the conditions required for implementing a dictatorial social choice function are also necessary for any nontrivial implementation. These conditions demand ex post optimality for the dictator, meaning her decision remains optimal regardless of the other agents' types. Making such an ex post optimal decision involves selecting the best alternative from a choice set Y. If Y contains more than one alternative, the decision—and thus the social choice function—becomes nontrivial.

From Proposition 4.1, it follows that if no agent can unambiguously compare any two distinct alternatives, then no one has a non-singleton decidable set, making the implementation of nontrivial social choice functions impossible:

Corollary 4.1. *A nontrivial implementable social choice function does not exist if, for every pair of distinct alternatives $x, y \in X$, every agent i has some type $\theta_i \in \Theta_i$ with which she does not unambiguously prefer x to y.*

Proposition 4.1, particularly Corollary 4.1, highlights that it is *ordinal* preference interdependence that obstructs implementation when ex post optimality is required. This finding aligns with the negative results in Feng, Niemeyer, and Wu (2023), which examines ex post implementation. In contrast, Feng and Wu (2023) shows that many nontrivial social choice functions can be implemented in a wide range of social choice problems if only interim strategy-proofness is required, without the necessity for ex post optimality.

Next, we apply the findings to two voting models, one with two alternatives and another with three, to further investigate their implications.

Binary voting

The agents must collectively decide whether to launch a safe project, which has a value of 0 for everyone, or a risky project, whose value is $v_i(\theta)$ for each agent i. The agents are willing to consider stochastic decisions, i.e. lotteries over the two projects, as well as deterministic choices. Agent i's preferences over these lotteries are represented by the utility function

$$u_i(p;\theta) = pv_i(\theta)$$

where p denotes the probability of launching the risky project.

For two lotteries p and p' where $p \neq p'$, agent i of type θ_i unambiguously prefers p to p' if and only if $(p - p')v_i(\theta_i, \theta_{-i})$ has the same sign for any $\theta_{-i} \in \Theta_{-i}$. This means that agent i can, by only using her own type, determine whether the risky project is either definitely better or definitely worse than the safe project. Therefore, by Corollary 4.1, as long as every agent has a type under which she is uncertain about this ordinal ranking, only trivial social choice functions are implementable.

Ternary voting

Now, let's increase the number of projects to three, which are the safe project (now denoted as s), a risky project a and a risky project b. The safe project continues to have a value of 0 for all agents, while the values of the risky projects a and b for agent i are given by $v_i^a(\theta)$ and $v_i^b(\theta)$, respectively. For simplicity, we will assume that the value of each risky project is never 0.

Agent i's preferences over lotteries involving these projects are represented by the utility function:

$$u_i(\alpha, \beta; \theta) = \alpha v_i^a(\theta) + \beta v_i^b(\theta).$$

Here, α and β respectively denote the probabilities with which risky projects a and b are launched.

By Corollary 4.1, nontrivial implementation is not possible if, for every agent i, there exists some type θ_i such that no pair of distinct lotteries can be compared unambiguously. In the Appendix[18], we show that this is the case if for every agent i, there is some type θ_i such that

$$[L_i^+(\theta_i), \quad U_i^+(\theta_i)] \cap [L_i^-(\theta_i), \quad U_i^-(\theta_i)] \neq \emptyset,$$

where

$$L_i^+(\theta_i) = \min_{\theta_{-i}} \left\{ \frac{v_i^a(\theta_i, \theta_{-i})}{v_i^b(\theta_i, \theta_{-i})} : v_i^b(\theta_i, \theta_{-i}) > 0 \right\}, \qquad U_i^+(\theta_i) = \max_{\theta_{-i}} \left\{ \frac{v_i^a(\theta_i, \theta_{-i})}{v_i^b(\theta_i, \theta_{-i})} : v_i^b(\theta_i, \theta_{-i}) > 0 \right\}$$

$$L_i^-(\theta_i) = \min_{\theta_{-i}} \left\{ \frac{v_i^a(\theta_i, \theta_{-i})}{v_i^b(\theta_i, \theta_{-i})} : v_i^b(\theta_i, \theta_{-i}) < 0 \right\}, \qquad U_i^-(\theta_i) = \max_{\theta_{-i}} \left\{ \frac{v_i^a(\theta_i, \theta_{-i})}{v_i^b(\theta_i, \theta_{-i})} : v_i^b(\theta_i, \theta_{-i}) < 0 \right\}.$$

This observation implies that nontrivial implementation is impossible if agents cannot consistently use their private information to effectively refine the potential ex post preference rankings among the three projects. For instance, this would be the case if every agent i has a type θ_i given which any of the following ex post preference rankings can occur:

$$a, b \succ_i s; \quad s \succ_i a, b; \quad a \succ_i s \succ_i b; \quad b \succ_i s \succ_i a.$$

Notice that in this instance, agent i of type θ_i cannot determine the ordinal ranking of any pair of projects. According to Corollary 4.1, this makes nontrivial implementation impossible if only *deterministic* social choice functions are considered. However, when lotteries are allowed, it is possible to construct examples[19] where, despite agents being unable to unambiguously compare any pair of distinct projects, some agents can unambiguously compare distinct lotteries over those projects. In such cases, nontrivial implementable social choice functions can indeed exist.

Appendix

Proof of Proposition 3.1

(This proof relies on the analysis in the Generalized One-dimensional Environments section. The reader is recommended to read that section before reading this proof.)

---

[18] Proposition 4.2

[19] See Appendix for an example.

Fix any agent i. For any two vectors of winning probabilities $p = (p_1, ..., p_n)$ and $p' = (p'_1, ..., p'_n)$, define $p \succsim_i^X p'$ if $p_i \geq p'_i$. For any two types of agent i, θ_i, θ'_i, define $\theta_i \succsim \theta'_i$ if θ_i is weakly higher than θ'_i. It follows that $\theta_i \sim_i^\Theta \theta'_i$ if θ_i and θ'_i are entangled, or $\theta_i \succ_i^\Theta \theta'_i$ if θ_i is higher than θ'_i.

We wish to show that the single-good auction environment manifests G1D in $(\succsim_i^X, \succsim_i^\Theta)_{i=1,...,n}$. The completeness and transitivity of $\succsim_i^X$ are evident. The completeness of $\succsim_i^\Theta$ follow from the fact that the "weakly higher" relation is complete, which can be verified straightforwardly.

Next, we show that $\succsim_i^\Theta$ is transitive. It is clear that $\sim_i^\Theta$ is an equivalence relation. Let $\mathcal{P}$ denote the partitioning of Θ_i induced by $\sim_i^\Theta$; that is, for any $P \in \mathcal{P}$, $\theta_i \in P$ and $\theta'_i \in P$ if and only if $\theta_i \sim_i^\Theta \theta'_i$. For each $P \in \mathcal{P}$, define $[P] := \cup_{\theta_i \in P}[\theta_i]$. Note that for any $P \neq P'$, we have $[P] \not\bowtie [P']$, for otherwise there would exist some $\theta_i \in P$ and $\theta'_i \in P'$ such that $[\theta_i] \bowtie [\theta'_i]$, or equivalently $\theta_i \sim_i^\Theta \theta'_i$, leading to a contradiction. Now, consider any θ_i, θ'_i , θ''_i such that $\theta_i \succsim_i^\Theta \theta'_i$ and $\theta'_i \succsim_i^\Theta \theta''_i$. and let P, P', P'' be the equivalence classes in $\mathcal{P}$ they respectively belong to. The condition $\theta_i \succsim_i^\Theta \theta'_i$ implies that either $P = P'$, or $[\theta'_i] \to [\theta_i]$, which implies $[P'] \to [P]$. Similarly, for $\theta'_i \succsim_i^\Theta \theta''_i$, we have either $P' = P''$ or $[P''] \to [P']$. There are four possible cases:

1. If $P = P' = P''$, then $\theta_i \sim_i^\Theta \theta''_i$, leading to $\theta_i \succsim_i^\Theta \theta''_i$.
2. If $P = P'$ and $[P''] \to [P']$, then $[P''] \to [P]$, implying $[\theta''_i] \to [\theta_i]$ and thus $\theta_i \succ_i^\Theta \theta''_i$, which gives $\theta_i \succsim_i^\Theta \theta''_i$.
3. Similarly, $\theta_i \succsim_i^\Theta \theta''_i$ holds if $[P'] \to [P]$ and $P' = P''$.
4. If $[P'] \to [P]$ and $[P''] \to [P']$, then clearly $[P''] \to [P]$, leading to $\theta_i \succsim_i^\Theta \theta''_i$.

Thus, we conclude that $\succsim_i^\Theta$ is transitive.

Next, we show that the environment manifests G1D in $(\succsim_i^X, \succsim_i^\Theta)_{i=1:n}$. Note that for any θ_i and $p, p' \in X$,

$$[\theta_i]^{pp'} = \begin{cases} \{(p_i - p'_i)\min_{\theta_{-i}} v_i(\theta_i, \theta_{-i}), (p_i - p'_i)\max_{\theta_{-i}} v_i(\theta_i, \theta_{-i})\} & \text{if } p_i \geq p'_i \\ \{(p_i - p'_i)\max_{\theta_{-i}} v_i(\theta_i, \theta_{-i}), (p_i - p'_i)\min_{\theta_{-i}} v_i(\theta_i, \theta_{-i})\} & \text{if } p_i < p'_i. \end{cases}$$

It is straightforward to verify that for any $\theta_i, \theta'_i \in \Theta_i$ and $p, p' \in X$:

- If $[\theta_i] \bowtie [\theta'_i]$, then $[\theta_i]^{pp'} \bowtie [\theta'_i]^{pp'}$ if $p_i \neq p'_i$;
- If $[\theta'_i] \to [\theta'_i]$, then $[\theta'_i]^{pp'} \to [\theta_i]^{pp'}$ if $p_i > p'_i$;
- $[\theta_i]^{pp'} = \{0\}$ if $p_i = p'_i$.

To verify Condition (1) in the definition of G1D, suppose $p \succ_i^X p'$ and $\theta_i \succ_i^\Theta \theta'_i$. Then, we have $p_i > p'_i$ and $[\theta'_i] \to [\theta_i]$. Hence, $[\theta'_i]^{pp'} \to [\theta_i]^{pp'}$.

To verify Condition (2), suppose $p \sim_i^X p'$. Then $p_i = p'_i$, and thus $[\theta_i]^{pp'} = \{0\}$.

To verify Condition (3), suppose $\theta_i \sim_i^\Theta \theta'_i$ and $\theta_i \neq \theta'_i$. There exists a sequence of agent i's types $\theta_i^1, ..., \theta_i^k$ such that $\theta_i^1 = \theta_i$, $\theta_i^k = \theta'_i$, and $[\theta_i] \bowtie [\theta'_i]$ for $\kappa = 1, ..., k-1$. This implies that $[\theta_i^\kappa]^{pp'} \sim [\theta_i^{\kappa+1}]^{pp'}$ if $p \not\sim_i^X p'$ (or equivalently $p_i \neq p'_i$). Moreover, $[\theta_i]^{pp'} \cup [\theta'_i]^{pp'}$ cannot be a singleton for any $p \not\sim_i^X p'$ because we have assumed that $v_i(\theta_i, \theta_{-i})$ is non-constant with respect to θ_{-i}, which implies that $[\theta_i]^{pp'}$ is not a singleton.

We have now established that the single-good auction environment manifests G1D in $(\succsim_i^X, \succsim_i^\Theta)_{i=1,\ldots,n}$. Pick any two types θ_i, θ_i' of agent i, and any report profile m_{-i} from the other agents. If θ_i is higher than θ_i', then $\theta_i \succ_i^\Theta \theta_i'$, and by Proposition 3.3, we have $q(\theta_i, m_{-i}) \succsim_i^X q(\theta_i', m_{-i})$, which implies that $q_i(\theta_i, m_{-i}) \geq q(\theta_i', m_{-i})$. If θ_i and θ_i' are entangled, then $\theta_i \sim_i^\Theta \theta_i'$, and by Proposition 3.3 we have $q(\theta_i, m_{-i}) \sim_i^X q(\theta_i', m_{-i})$, which implies that $q_i(\theta_i, m_{-i}) = q_i(\theta_i', m_{-i})$.

Proof of Proposition 3.2

(This proof relies on the analysis in the Generalized One-dimensional Environments section. The reader is recommended to read that section before reading this proof.)

Let $(\succsim_i^X, \succsim_i^\Theta)_{i=1:n}$ be the rankings of alternatives and types defined for the single-auction environment in the proof of Proposition 3.1. We have shown in the proof of Proposition 3.1 that the environment manifests G1D in them.

“Only if”: Suppose q^E is implementable. Pick any $\theta \in \Theta$. Suppose $i^*(\theta) = i$. Consider any θ_i' that is weakly higher than θ_i (and hence $\theta_i' \succsim_i^\Theta \theta_i$). By Proposition 3.3, we have $q_i^E(\theta_i', \theta_{-i}) \geq q_i^E(\theta) = 1$, which implies $i = i^*(\theta_i', \theta_{-i})$. Therefore the environment satisfies strong single-crossing.

“If”: Suppose the environment satisfies strong single-crossing. Consider the following transfer functions τ_i, $i = 1, \ldots, n$:

$$\tau_i(\theta) = \begin{cases} 0 & \text{if } i \neq i^*(\theta) \\ \max_{\{\tilde{\theta} \in \Theta : i \neq i^*(\tilde{\theta}_i, \theta_{-i})\}} v_i(\tilde{\theta}) & \text{if } i = i^*(\theta). \end{cases}$$

Fix any agent i, type profile θ, and reports m_{-i} from the other agents. First, consider the case where $i \neq i^*(\theta_i, m_{-i})$. To verify that agent i has no incentive to misreport under type profile θ while the other agents report m_{-i}, we only need to check reporting some θ_i' such that $i = i^*(\theta_i', m_{-i})$ is unprofitable, because only such a deviation leads to a change in her probability of winning and transfer payment. If agent i reports her true type θ_i, her utility is 0. If she instead deviates to reporting such a θ_i', she would win the good with certainty and receive a payoff of $v_i(\theta)$, while having to pay $\max_{\{\tilde{\theta} \in \Theta : i \neq i^*(\tilde{\theta}_i, m_{-i})\}} v_i(\tilde{\theta})$, which is no less than $v_i(\theta)$ because $i \neq i^*(\theta_i, m_{-i})$. Therefore, such a deviation is unprofitable.

Next, consider the case where $i = i^*(\theta_i, m_{-i})$. It is sufficient to verify that a deviation to reporting θ_i' such that $i \neq i^*(\theta_i', m_{-i})$ is unprofitable. Agent i's utility from reporting such a θ_i' is 0, whereas by reporting her true type, she wins the good with certainty, receives a payoff of $v_i(\theta)$, and pays $\max_{\{\tilde{\theta} \in \Theta : i \neq i^*(\tilde{\theta}_i, m_{-i})\}} v_i(\tilde{\theta})$. Since $i = i^*(\theta_i, m_{-i})$, by strong single-crossing we have $\theta_i \succsim_i^\Theta \tilde{\theta}_i$ for any $\tilde{\theta}_i$ where $i \neq i^*(\tilde{\theta}_i, m_{-i})$, and hence:

$$v_i(\theta) \geq \min_{\tilde{\theta}_{-i}} v_i(\theta_i, \tilde{\theta}_{-i}) \geq \max_{\tilde{\theta}_{-i}} v_i(\tilde{\theta}_i, \tilde{\theta}_{-i}) \qquad \forall \tilde{\theta}_i \quad s.t. \quad i \neq i^*(\tilde{\theta}_i, m_{-i}),$$

and it follows that $v_i(\theta) \geq \max_{\{\tilde{\theta} \in \Theta : i \neq i^*(\tilde{\theta}_i, m_{-i})\}} v_i(\tilde{\theta}) = \tau_i(\theta_i, m_{-i})$. Thus, deviating to report θ_i' is unprofitable.

Proof of Proposition 3.3

"If": Suppose q satisfies monotonicity and screenability as stipulated in the proposition. Fix any agent i and any report profile m_{-i} from the other agents. Let $\hat{X}$ denote the range of $q(\cdot, m_{-i})$. For simplicity, assume that for any distinct $x, y \in \hat{X}$, we have either $x \succ_i^X y$ or $y \succ_i^X x$. The current proof can be easily adapted to allow for the case where $x \sim_i^X y$ by treating such x and y as if they are the same alternative.

Index the alternatives in $\hat{X}$ as $x^1, ..., x^l$ such that $x^l \succ_i^X x^{l-1} \succ_i^X ... \succ_i^X x^1$. For each $k = 1, ..., l$, denote $\Theta_i^k := \{\theta_i : q(\theta_i, m_{-i}) = x^k\}$, and for $k = 2, ..., l$, define

$$\tau^k := \min_{\theta_i \in \Theta_i^k} \min_{\theta_{-i}} u_i^{x^k x^{k-1}}(\theta_i, \theta_{-i}).$$

Due to monotonicity and screenability, if $\theta_i \in \Theta_i^k$ and $\theta_i' \in \Theta_i^{k'}$ with $k > k'$ (or equivalently $q(\theta_i, m_{-i}) \succ_i^X q(\theta_i', m_{-i})$), then $\theta_i \succ_i^\Theta \theta_i'$, which implies that $[\theta_i']^{x^k x^{k-1}} \to [\theta_i]^{x^k x^{k-1}}$. Thus, for any $k = 2, ..., l$, we have

$$\tau^k \geq \max_{\theta_{-i}} u_i^{x^k x^{k-1}}(\theta_i', \theta_{-i}) \qquad \forall \theta_i' \in \Theta_i^{k'} \text{ where } k' < k \tag{I1}$$

$$\tau^k \leq \min_{\theta_{-i}} u_i^{x^k x^{k-1}}(\theta_i', \theta_{-i}) \qquad \forall \theta_i' \in \Theta_i^{k'} \text{ where } k' > k. \tag{I2}$$

Given m_{-i}, construct the transfers for agent i as follows:

$$\tau_i(\theta_i, m_{-i}) = \begin{cases} 0 & \text{if } \theta_i \in \Theta_i^1 \\ \sum_{\kappa=2}^k \tau^\kappa & \text{if } \theta_i \in \Theta_i^k \text{ where } k \geq 2. \end{cases}$$

Now we wish to show that if the other agents report m_{-i}, truthful reporting is optimal for agent i regardless of the actual types of the other agents. Suppose agent i's type is θ_i. Thus $\theta_i \in \Theta_i^k$ for some $k \in \{1, ..., l\}$. Let the types of the other agents be θ_{-i}.

Observe that agent i has no incentive to deviate to reporting any $\theta_i' \in \Theta_i^k$, because it changes neither the chosen alternative nor her transfer payment. If she deviates to reporting $\theta_i' \in \Theta_i^{k'}$ where $k' > k$, the resulting change in her utility is

$$\begin{aligned} u_i^{x^{k'} x^k}(\theta_i, \theta_{-i}) - \sum_{\kappa=k+1}^{k'} \tau^\kappa &= \sum_{\kappa=k+1}^{k'} u_i^{x^\kappa x^{\kappa-1}}(\theta_i, \theta_{-i}) - \sum_{\kappa=k+1}^{k'} \tau^\kappa \\ &= \sum_{\kappa=k+1}^{k'} \left[u_i^{x^\kappa x^{\kappa-1}}(\theta_i, \theta_{-i}) - \tau^\kappa \right] \leq 0 \end{aligned}$$

where the final inequality follows from (I1). Deviating to report any $\theta_i' \in \Theta_i^{k'}$ where $k' < k$ can be shown to be similarly unprofitable (by invoking inequality (I2)). Since m_{-i} is arbitrarily chosen, it follows that q is implementable.

"Only if": Suppose q is implementable with supporting transfer functions $\tau_1, ..., \tau_n$. Fix any agent i and any report profile m_{-i} from the other agents. Pick a pair of agent i's types θ_i and θ_i'. Denote $x := q(\theta_i, m_{-i})$ and $y := q(\theta_i', m_{-i})$. Incentive compatibility for agent i of type θ_i requires $u_i(x; \theta_i, \theta_{-i}) - \tau_i(\theta_i, m_{-i}) \geq u_i(y; \theta_i, \theta_{-i}) - \tau_i(\theta_i', m_{-i}) \quad \forall \theta_{-i} \in \Theta_{-i}$, which implies $\min_{\theta_{-i}} u_i^{xy}(\theta_i, \theta_{-i}) \geq \tau_i(\theta_i, m_{-i}) - \tau_i(\theta_i', m_{-i})$. Similarly, incentive compatibility for agent i of type θ_i' implies $\tau_i(\theta_i, m_{-i}) - \tau_i(\theta_i', m_{-i}) \geq \max_{\theta_{-i}} u_i^{xy}(\theta_i', \theta_{-i})$. Thus we have

$$\min_{\theta_{-i}} u_i^{xy}(\theta_i, \theta_{-i}) \geq \max_{\theta_{-i}} u_i^{xy}(\theta_i', \theta_{-i}) \quad \Longleftrightarrow \quad \max_{\theta_{-i}} u_i^{yx}(\theta_i, \theta_{-i}) \leq \min_{\theta_{-i}} u_i^{yx}(\theta_i', \theta_{-i}). \tag{I3}$$

To show monotonicity, suppose $\theta_i \succ_i^\Theta \theta_i'$, and we wish to show that $x \succsim_i^X y$. Observe that $y \not\succ_i^X x$, for otherwise we would have $[\theta_i']^{yx} \rightarrow [\theta_i]^{yx}$, implying $\max_{\theta_{-i}} u_i^{yx}(\theta_i, \theta_{-i}) > \min_{\theta_{-i}} u_i^{yx}(\theta_i', \theta_{-i})$, which contradicts Inequality (I3). Thus we have $x \succsim_i^X y$.

To show screenability, suppose $\theta_i \sim_i^\Theta \theta_i'$, and we wish to show that $x \sim_i^X y$. Suppose, for the sake of contradiction, that $x \not\sim_i^X y$. Then Condition (3a) implies that there exists a sequence of types $\theta_i^1, ..., \theta_i^k$ such that $\theta_i^1 = \theta_i$, $\theta_i^k = \theta_i'$, and $[\theta_i^\kappa]^{xy} \bowtie [\theta_i^{\kappa+1}]^{xy}$ for $\kappa = 1, ..., k-1$. Note that for each $\kappa = 1, ..., k-1$, we have $\theta_i^\kappa \sim_i^\Theta \theta_i^{\kappa+1}$ (otherwise, by Condition (1), we would have either $[\theta_i^\kappa]^{xy} \rightarrow [\theta_i^{\kappa+1}]^{xy}$ or $[\theta_i^{\kappa+1}]^{xy} \rightarrow [\theta_i^\kappa]^{xy}$, since $x \not\sim_i^X y$). Fix any $\kappa \in \{1, ..., k-1\}$. Denote $w := q(\theta_i^\kappa, m_{-i})$ and $z := q(\theta_i^{\kappa+1}, m_{-i})$. We want to show that $w \sim_i^X z$. Suppose $w \not\sim_i^X z$, then by Condition (3a), $[\theta_i^{\kappa+1}]^{wz} \bowtie [\theta_i^\kappa]^{wz}$.

Now the incentive compatibility for agent i of type θ_i^κ to not deviate to reporting type $\theta_i^{\kappa+1}$ implies that $\min_{\theta_{-i}} u_i^{wz}(\theta_i^\kappa, \theta_{-i}) \geq \tau_i(\theta_i^\kappa, m_{-i}) - \tau_i(\theta_i^{\kappa+1}, m_{-i})$. Similarly, for agent i of type $\theta_i^{\kappa+1}$ to not deviate to reporting type θ_i^κ, we must have $\tau_i(\theta_i^\kappa, m_{-i}) - \tau_i(\theta_i^{\kappa+1}, m_{-i}) \geq \max_{\theta_{-i}} u_i^{wz}(\theta_i^{\kappa+1}, \theta_{-i})$. Thus we derive that

$$\min_{\theta_{-i}} u_i^{wz}(\theta_i^\kappa, \theta_{-i}) \geq \max_{\theta_{-i}} u_i^{wz}(\theta_i^{\kappa+1}, \theta_{-i}). \tag{I4}$$

Recall that there are three possibilities (see page 13) for $[\theta_i^{\kappa+1}]^{wz} \bowtie [\theta_i^\kappa]^{wz}$. Inequality (I4) leads to the third possibility: the two intervals coincide as degenerate intervals, which violates Condition (3b).

Thus we must have $w \sim_i^X z$. Since κ is arbitrarily chosen from $\{1, ..., k\}$, we conclude that $x = q(\theta_i^1, m_{-i}) \sim_i^X q(\theta_i^2, m_{-i}) \sim_i^X ... \sim_i^X q(\theta_i^k, m_{-i}) = y$, which contradicts the assumption that $x \not\sim_i^X y$. It follows that $x \sim_i^X y$.

Proof of Propositions 3.4

Given any agent i, any finite sequence of her types $\boldsymbol{\theta_i} = (\theta_i^1, ..., \theta_i^k)$, and any report profile m_{-i} from the other agents, define

$$G(\boldsymbol{\theta_i}, m_{-i}) := \sum_{\kappa=1}^{k-1} \max_{\theta_{-i}} \Big\{ u_i\Big(q(\theta_i^{\kappa+1}, m_{-i}); \theta_i^\kappa, \theta_{-i}\Big) - u_i\Big(q(\theta_i^\kappa, m_{-i}); \theta_i^\kappa, \theta_{-i}\Big) \Big\}.$$

"If": Suppose q satisfies GCM. Fix an agent i and a type $\tilde{\theta}_i \in \Theta_i$. For any $\theta_i \in \Theta_i$, let $S(\theta_i)$ denote the set of all finite sequences of types of i that starts with $\tilde{\theta}_i$ and terminates with θ_i.

For any $\theta_i \in \Theta_i$ and report profile $m_{-i} \in \Theta_{-i}$, define

$$V(\theta_i, m_{-i}) := \sup_{\boldsymbol{\theta_i} \in S(\theta_i)} G(\boldsymbol{\theta_i}, m_{-i}).$$

Note that GCM implies $V(\tilde{\theta}_i, m_{-i}) \leq 0$. Moreover, since $G(\boldsymbol{\theta_i}, m_{-i}) = 0$ for $\boldsymbol{\theta_i} = (\tilde{\theta}_i, \tilde{\theta}_i)$, we have $V(\tilde{\theta}_i, m_{-i}) = 0$ for any $m_{-i} \in \Theta_{-i}$.

Observe that

$$\begin{aligned} V(\tilde{\theta}_i, m_{-i}) \geq & \sup_{\boldsymbol{\theta_i} \in S(\tilde{\theta}_i), \theta_i^{k-1} = \theta_i} G(\boldsymbol{\theta_i}, m_{-i}) \\ &= V(\theta_i, m_{-i}) + \max_{\theta_{-i}} \Big\{ u_i\Big(q(\tilde{\theta}_i, m_{-i}), \theta_i, \theta_{-i}\Big) - u_i\Big(q(\theta_i, m_{-i}); \theta_i, \theta_{-i}\Big) \Big\}. \end{aligned}$$

It follows that $V(\theta_i, m_{-i}) \leq -\max_{\theta_{-i}} \Big\{ u_i\Big(q(\tilde{\theta}_i, m_{-i}), \theta_i, \theta_{-i}\Big) - u_i\Big(q(\theta_i, m_{-i}); \theta_i, \theta_{-i}\Big) \Big\}$. Thus $V(\cdot, m_{-i})$ is well defined.

Consider the transfer function $\tau_i(\theta_i, m_{-i}) := V(\theta_i, m_{-i})$. Observe that for any $\theta_i, \theta_i' \in \Theta_i$, $\theta_{-i}, m_{-i} \in \Theta_{-i}$,

$$\begin{aligned} & \Big[u_i\Big(q(\theta_i, m_{-i}); \theta_i, \theta_{-i}\Big) - \tau_i(\theta_i, m_{-i})\Big] - \Big[u_i\Big(q(\theta_i', m_{-i}); \theta_i, \theta_{-i}\Big) - \tau_i(\theta_i', m_{-i})\Big] \\ =& V(\theta_i', m_{-i}) - \Big[V(\theta_i, m_{-i}) + u_i\Big(q(\theta_i', m_{-i}); \theta_i, \theta_{-i}\Big) - u_i\Big(q(\theta_i, m_{-i}); \theta_i, \theta_{-i}\Big)\Big] \\ \geq& V(\theta_i', m_{-i}) - \Big[V(\theta_i, m_{-i}) + \max_{\hat{\theta}_{-i}} \Big\{ u_i\Big(q(\theta_i', m_{-i}); \theta_i, \hat{\theta}_{-i}\Big) - u_i\Big(q(\theta_i, m_{-i}); \theta_i, \hat{\theta}_{-i}\Big) \Big\}\Big] \\ =& V(\theta_i', m_{-i}) - \sup_{\boldsymbol{\theta_i} \in S(\theta_i'), \theta_i^{k-1} = \theta_i} G(\boldsymbol{\theta_i}, m_{-i}) \geq 0. \end{aligned}$$

It follows that q is implementable.

“Only if”: Suppose q is implementable with supporting transfer functions $\tau_1, ..., \tau_n$. Fix any agent i and any sequence $(\theta_i^1, ..., \theta_i^k)$ of types of agent i where $\theta_i^k = \theta_i^1$. Incentive compatibility implies that for any $\kappa = 1, ..., k-1$ and any report profile $m_{-i} \in \Theta_{-i}$,

$$\forall \theta_{-i} \in \Theta_{-i} : u_i\Big(q(\theta_i^{\kappa+1}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) - \tau_i(\theta_i^{\kappa+1}, m_{-i}) \leq u_i\Big(q(\theta_i^{\kappa}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) - \tau_i(\theta_i^{\kappa}, m_{-i}),$$

which in turn implies that

$$\max_{\theta_{-i}} \Big\{ u_i\Big(q(\theta_i^{\kappa+1}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) - u_i\Big(q(\theta_i^{\kappa}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) \Big\} \leq \tau_i(\theta_i^{\kappa+1}, m_{-i}) - \tau_i(\theta_i^{\kappa}, m_{-i}).$$

Denote the above inequality as I_κ. If we sum up $I_1, I_2, ..., I_{k-1}$, it is clear that all the terms on the right hand side will cancel out, and we are left with

$$\sum_{\kappa=1}^{k-1} \max_{\theta_{-i}} \Big\{ u_i\Big(q(\theta_i^{\kappa+1}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) - u_i\Big(q(\theta_i^{\kappa}, m_{-i}); \theta_i^{\kappa}, \theta_{-i}\Big) \Big\} \leq 0.$$

It follows that q satisfies GCM.

Proof of Propositions 3.5

Here we prove that reshuffling-efficiency is equivalent to GCM, thereby establishing Proposition 3.5 based on Proposition 3.4. We inherit the definition of $G(\boldsymbol{\theta_i}, m_{-i})$ from the proof of Proposition 3.4.

Suppose q satisfies GCM. Fix any agent i and any report profile $m_{-i} \in \Theta_{-i}$. Let μ be the type-alternative matching for agent i induced by q given m_{-i}. Choose any reshuffling μ' of μ, and let π be the permutation used to obtain μ' from μ, i.e. $\mu'(\theta_i) = \mu(\pi(\theta_i))$ for any $\theta_i \in \Theta_i$.

Since π is a permutation of Θ_i, it induces a permutation group on Θ_i, implying that Θ_i can be decomposed into disjoint permutation cycles,[20] i.e. there exists a partition P of Θ_i such that for any $p \in P$ the types in p can be ordered as $\theta_i^1, ..., \theta_i^{k-1}$ (where $k-1$ denotes the number of types in p) satisfying (1) $\theta_i^1 = \pi(\theta_i^{k-1})$, and (2) $\theta_i^{\kappa+1} = \pi(\theta_i^{\kappa})$ for every $\kappa = 1, ..., k-2$. Pick any $p \in P$ and order the types in it as $\theta_i^1, ..., \theta_i^{k-1}$ as described in the previous sentence. Define $\boldsymbol{\theta_i} = (\theta_i^1, ..., \theta_i^{k-1}, \theta_i^1)$ and $W(p) := G(\boldsymbol{\theta_i})$.

Note that $\boldsymbol{\theta_i}$ is a sequence of types where the initial and terminal terms are the same. Therefore $W(p) = G(\boldsymbol{\theta_i}, m_{-i}) \leq 0$ by GCM, and it follows that $\sum_{p \in P} W(p) \leq 0$. We observe that

$$D_i(\mu', \mu) = \sum_{\theta_i \in \Theta_i} \max_{\theta_{-i}} \left\{ u_i\Big(\underbrace{q(\pi(\theta_i), m_{-i})}_{\mu'(\theta_i)}; \theta_i, \theta_{-i}\Big) - u_i\Big(\underbrace{q(\theta_i, m_{-i})}_{\mu(\theta_i)}; \theta_i, \theta_{-i}\Big) \right\} = \sum_{p \in P} W(p) \leq 0.$$

Thus μ is reshuffling-efficient.

Now suppose that for any agent i and any report profile m_{-i} from the other agents, the type-alternative matching μ for agent i induced by q given by m_{-i} is reshuffling-efficient. However, suppose GCM is not satisfied. Then, for some m_{-i}, there exists at least one sequence $\boldsymbol{\theta_i} = (\theta_i^1, ..., \theta_i^k)$ where $\theta_i^1 = \theta_i^k$ such that $G(\boldsymbol{\theta_i}, m_{-i}) > 0$. Let $\boldsymbol{\hat{\theta}_i} = (\hat{\theta}_i^1, ..., \hat{\theta}_i^k)$ be the *shortest* of such sequences given m_{-i}.

First, consider the case where $\boldsymbol{\hat{\theta}_i}$ forms a cycle, i.e. there are no repetitive terms except for the coincidence of the initial and the terminal terms. Consider the function $\pi : \Theta_i \to \Theta_i$ such that

$$\pi(\theta_i) = \begin{cases} \hat{\theta}_i^{\kappa+1} & \text{if } \theta_i = \hat{\theta}_i^{\kappa} \text{ for some } \kappa = 1, ..., k-1 \\ \theta_i & \text{otherwise.} \end{cases}$$

It is straightforward to verify that π is a permutation, as each type appearing in the sequence $\boldsymbol{\hat{\theta}_i}$ is mapped to the next type in the sequence (where $\hat{\theta}_i^k$ is ignored because it has already appeared as $\hat{\theta}_i^1$), while types not in the sequence map to themselves. Let μ'

---

[20] See, for example, Theorem 9.8 of Fraleigh (2002), pp. 89.

be the reshuffling obtained from π, i.e. $\mu'(\theta_i) = \mu(\pi(\theta_i))$. Thus we have

$$\begin{aligned} D_i(\mu',\mu) = &\sum_{\theta_i \in \{\hat{\theta}_i^1,\ldots,\hat{\theta}_i^{k-1}\}} \max_{\theta_{-i}} \Big\{ u_i\Big(\mu'(\theta_i);\theta_i,\theta_{-i}\Big) - u_i\Big(\mu(\theta_i);\theta_i,\theta_{-i}\Big)\Big\} \\ &+ \sum_{\theta_i \notin \{\hat{\theta}_i^1,\ldots,\hat{\theta}_i^{k-1}\}} \max_{\theta_{-i}} \Big\{ u_i\Big(\underbrace{\mu'(\theta_i)}_{=\mu(\theta_i)};\theta_i,\theta_{-i}\Big) - u_i\Big(\mu(\theta_i);\theta_i,\theta_{-i}\Big)\Big\} = G(\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}, m_{-i}) > 0, \end{aligned}$$

which contradicts the assumption that μ is reshuffling-efficient.

Now consider the case where $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$ is not a cycle. In this case, there exist a, b where $1 \leq a < b \leq k$ and $a = 1$ and $b = k$ do not hold simultaneously, such that $(\hat{\theta}_i^a, ..., \hat{\theta}_i^b)$ forms a cycle. By the argument in the previous paragraph, if $\sum_{i=a}^{b-1} \max_{\theta_{-i}} \Big\{ u_i\Big(q(\hat{\theta}_i^{\kappa+1}, m_{-i}); \hat{\theta}_i^{\kappa}, \theta_{-i}\Big) - u_i\Big(q(\hat{\theta}_i^{\kappa}, m_{-i}); \hat{\theta}_i^{\kappa}, \theta_{-i}\Big)\Big\}$ is strictly positive, then μ is not reshuffling-efficient. Therefore, this summation must be non-positive. Consider another sequence $\tilde{\boldsymbol{\theta}}_{\boldsymbol{i}} = (\tilde{\theta}_i^1, ..., \tilde{\theta}_i^{k'})$ obtained from $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$ by removing $(\hat{\theta}_i^{a+1}, ..., \hat{\theta}_i^b)$, i.e. $(\tilde{\theta}_i^1, ..., \tilde{\theta}_i^{k'}) = (\hat{\theta}_i^1, ..., \hat{\theta}_i^a, \hat{\theta}_i^{b+1}, ..., \hat{\theta}_i^k)$. It is straightforward to verify that $\tilde{\theta}_i^1 = \tilde{\theta}_i^{k'}$. Observe that

$$G(\tilde{\boldsymbol{\theta}}_{\boldsymbol{i}}, m_{-i}) = \underbrace{G(\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}, m_{-i})}_{>0} - \underbrace{\sum_{\kappa=a}^{b-1} \max_{\theta_{-i}} \Big\{ u_i\Big(q(\hat{\theta}_i^{\kappa+1}, m_{-i}); \hat{\theta}_i^{\kappa}, \theta_{-i}\Big) - u_i\Big(q(\hat{\theta}_i^{\kappa}, m_{-i}); \hat{\theta}_i^{\kappa}, \theta_{-i}\Big)\Big\}}_{\leq 0} > 0.$$

However, this contradicts the choice of $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$, which is the shortest sequence of agent i's types with identical initial and terminal terms under which $G(\cdot, m_{-i})$ is positive, because $\tilde{\boldsymbol{\theta}}_{\boldsymbol{i}}$ is by construction strictly shorter than $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$. This concludes the proof.

Proof of Proposition 4.1

"If": Suppose some non-singleton $Y \subset X$ is decidable for agent i. Let c_i be the associated decision function. Consider the social choice function $q(\theta_i, m_{-i}) := c_i(\theta_i)\ \forall m_{-i} \in \Theta_{-i}$. Since the range of c_i is Y, which is non-singleton, q is nontrivial. To show that q is implementable, first observe that q does not depend on the report from any agent $j \neq i$, and hence agent j has no incentive to misreport her type. As for agent i, by reporting her true type θ_i, she can ensure that $c_i(\theta_i)$ is the chosen alternative, which she unambiguously prefers to $c_i(\theta_i')$, which is what she would induce should she deviate to reporting θ_i'.

"Only if": Suppose there exists a nontrivial, implementable social choice function q. Since q is nontrivial, there must exist an agent i such that $q(\theta_i, m_{-i}) \neq q(\theta_i', m_{-i})$ for some $\theta_i, \theta_i' \in \Theta_i, m_{-i} \in \Theta_{-i}$. Let $Y := \{q(\theta_i, m_{-i}) : \theta_i \in \Theta_i\}$ and $c_i(\theta_i) := q(\theta_i, m_{-i})$. By construction, $|Y| > 1$ and c_i is surjective when its codomain is restricted to Y. By Lemma 4.1, agent i of type θ_i unambiguously prefers $c_i(\theta_i)$ to $c_i(\theta_i')$ for any $\theta_i' \in \Theta_i$. Hence Y is decidable for agent i.

Derivation of the result in ternary voting

Proposition 4.2. *There does not exist a nontrivial implementable social choice function if for every agent $i \in I$, there exists $\theta_i \in \Theta_i$ such that $[L_i^+(\theta_i), U_i^+(\theta_i)] \cap [L_i^-(\theta_i), U_i^-(\theta_i)] \neq \emptyset$.*

Proof. For any agent i, type profile $\theta = (\theta_i, \theta_{-i})$ and lottery (α, β), let $L_i(\alpha, \beta; \theta)$ denote the lower contour set of (α, β) under agent i's preferences conditional on θ, i.e. $L_i(\alpha, \beta; \theta) = \{(\alpha', \beta') : \alpha' v_i^a(\theta) + \beta' v_i^b(\theta) \leq \alpha v_i^a(\theta) + \beta v_i^b(\theta)\}$. The set of lotteries that agent i of type θ_i finds unambiguously no better than (α, β) is thus

$$\bigcap_{\theta_{-i}} L_i(\alpha, \beta; \theta_i, \theta_{-i}) = \underbrace{\left(\bigcap_{\theta_{-i}: v_i^b(\theta_i, \theta_{-i}) > 0} L_i(\alpha, \beta; \theta_i, \theta_{-i}) \right)}_{S} \bigcap \underbrace{\left(\bigcap_{\theta_{-i}: v_i^b(\theta_i, \theta_{-i}) < 0} L_i(\alpha, \beta; \theta_i, \theta_{-i}) \right)}_{T}.$$

Observe that $S = \big\{(\alpha', \beta') : \beta' \leq \max\{L_i^+(\theta_i)(\alpha - \alpha'), U_i^+(\theta_i)(\alpha - \alpha')\} + \beta\big\}$ and $T = \big\{(\alpha', \beta') : \beta' \geq \min\{L_i^-(\theta_i)(\alpha - \alpha'), U_i^-(\theta_i)(\alpha - \alpha')\} + \beta\big\}$. Note that $S \cap T = \{(\alpha, \beta)\}$ if and only if $[L_i^+(\theta_i), U_i^+(\theta_i)] \cap [L_i^-(\theta_i), U_i^-(\theta_i)] \neq \emptyset$. The proposition then follows from Corollary 4.1. ☐

Ternary voting (continued)

There are two agents, 1 and 2, each of whom may be of either the H type or the L type. The values of the projects are given in the following table.

		θ_{-i}					
		H			L		
θ_i	H	$a:2$	$s:0$	$b:-1$	$a:-1$	$s:0$	$b:2$
	L	$a:1$	$s:0$	$b:-2$	$a:-2$	$s:0$	$b:1$

Clearly, neither agent can unambiguously compare any pair of distinct projects based solely on her private information.

Now let's consider two lotteries:

1. Risky project a and risky project b are each launched with a probability of 0.5;
2. The safe project is launched with certainty.

It is straightforward to verify that an agent unambiguously prefers lottery (1) to lottery (2) if her type is H, and vice versa if her type is L. Therefore, the two lotteries form a non-singleton decidable set for either agent. This leads us to conclude, based on Proposition 4.1, that nontrivial implementation is indeed possible in this case.

References

Barberà, Salvador, Dolors Berga, and Bernardo Moreno (2019), "Domains admitting ex post incentive compatible and respectful mechanisms: A characterization for the two alternatives case." *Social Design: Essays in Memory of Leonid Hurwicz*, 295–306.

Barberà, Salvador, Dolors Berga, and Bernardo Moreno (2022), "Restricted environments and incentive compatibility in interdependent values models." *Games and Economic Behavior*, 131, 1–28.

Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813.

Bergemann, Dirk and Juuso Välimäki (2002), "Information acquisition and efficient mechanism design." *Econometrica*, 70, 1007–1033.

Bikhchandani, Sushil (2006), "Ex post implementation in environments with private goods." *Theoretical Economics*, 1.

Börgers, Tilman (2015), *An Introduction to the Theory of Mechanism Design*. Oxford University Press.

Brooks, Benjamin and Songzi Du (2024), "On the structure of informationally robust optimal mechanisms." *Econometrica*, 92, 1391–1438.

Che, Yeon-Koo, Jinwoo Kim, and Fuhito Kojima (2015), "Efficient assignment with interdependent values." *Journal of Economic Theory*, 158, 54–86.

Crémer, Jacques and Richard P McLean (1985), "Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent." *Econometrica*, 53, 345–361.

Dasgupta, Partha and Eric Maskin (2000), "Efficient auctions." *The Quarterly Journal of Economics*, 115, 341–388.

Dworczak, Piotr and Anthony Lee Zhang (2017), "Implementability, walrasian equilibria, and efficient matchings." *Economics Letters*, 153, 57–60.

Feng, Tangren, Axel Niemeyer, and Qinggong Wu (2023), "The limits of ex post implementation without transfers." *Theoretical Economics*, 18, 463–479.

Feng, Tangren and Qinggong Wu (2023), "Interim strategy-proofness." *Working paper*.

Fraleigh, John B. (2002), *A First Course in Abstract Algebra, 7th Edition*. Pearson.

Fujinaka, Yuji and Toshiji Miyakawa (2020), "Ex-post incentive compatible and individually rational assignments in housing markets with interdependent values." *Journal of Mathematical Economics*, 91, 157–164.

Jehiel, Philippe, Moritz Meyer-ter Vehn, Benny Moldovanu, and William R Zame (2006), "The limits of ex post implementation." *Econometrica*, 74, 585–610.

Jehiel, Philippe and Benny Moldovanu (2001), "Efficient design with interdependent valuations." *Econometrica*, 69, 1237–1259.

Kos, Nenad and Matthias Messner (2013), "Extremal incentive compatible transfers." *Journal of Economic Theory*, 148, 134–164.

Maskin, Eric (1992), *Auctions and Privatization*, 115–136. J.C.B. Mohr Publisher.

Myerson, Roger B. (1981), “Optimal auction design.” *Mathematics of Operations Research*, 6, 58–73.

Niemeyer, Axel and Justus Preusser (2024), “Optimal allocation with peer information.” *Working paper*.

Perry, Motty and Philip J. Reny (2002), “An efficient auction.” *Econometrica*, 70, 1199–1212.

Pourpouneh, Mohsen, Rasoul Ramezanian, and Arunava Sen (2020), “The marriage problem with interdependent preferences.” *International Game Theory Review (IGTR)*, 22, 1–16.

Rahman, David M. (2024), “Detecting profitable deviations.” *Journal of Mathematical Economics*, 111, 102946.

Rochet, Jean-Charles (1987), “A necessary and sufficient condition for rationalizability in a quasi-linear context.” *Journal of Mathematical Economics*, 16, 191–200.

Rockafellar, R. Tyrrell (1970), *Convex analysis*. Princeton Mathematical Series, Princeton University Press, Princeton, N. J.

Tang, Rui and Mu Zhang (2021), “Maxmin implementation.” *Journal of Economic Theory*, 194, 105250.

Williams, Steven R and Roy Radner (1988), “Informational externalities and the scope of efficient dominant strategy mechanisms.” *Working paper*.