Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Ex Post Strategy-Proofness

Tangren Feng and Qinggong Wu

Working Paper n. 724

This Version: February 5, 2026

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

Ex Post Strategy-Proofness*

Tangren Feng† and Qinggong Wu‡

February 5, 2026

Abstract

Ex post strategy-proofness (ESP) extends the concept of strategy-proofness to settings with interdependent values. ESP requires that truthful reporting is always optimal for every agent, regardless of the strategies *and* types of other agents. We characterize the class of social choice functions implementable by ESP mechanisms with transfers. For a broad class of environments, including single-good auctions, implementability is jointly characterized by two conditions: (1) *Monotonicity*, requiring that "higher" types induce "higher" alternatives; and (2) *Screenability*, requiring that types which are "entangled" must be treated identically, where entanglement becomes increasingly restrictive as types become denser or preference interdependence strengthens. For general settings, we establish that implementability is equivalent to a generalized cyclical monotonicity condition, or equivalently, to an optimality condition derived from an induced matching problem between types and alternatives.

Keywords: Mechanism design, interdependent values, strategy-proofness, dominant strategy, auction design, simplicity, strategic uncertainty

1 Introduction

In mechanism design, we often aim for strategy-proofness. By ensuring that every agent has a dominant strategy, strategy-proof mechanisms offer simplicity and robustness. However, the vast majority of existing results on strategy-proof mechanism design rely on the assumption of *private values*, where an agent's preferences depend only on her own private information. In contrast, strategy-proofness is far less understood in *interdependent values* settings, where an agent's preferences are influenced by the private information held by others. Yet, such interdependence is central to many economically significant scenarios. Prominent examples include common value auctions, where bidders refine their valuations based on competitors' signals, and jury or committee voting, where information about the "true state" is dispersed among 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.*

This paper addresses this gap. Adopting the standard (partial) implementation approach, we characterize the class of social choice functions that can be implemented by a strategy-proof mechanism in interdependent value settings with transfers. In doing so, our analysis 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 the definition of a dominant strategy within interdependent value settings. Under the conventional definition—typically applied to private value settings without preference interdependence—a strategy is dominant if it guarantees optimality regardless of the strategies chosen by other agents. Extending this definition to accommodate interdependent values, however, is not straightforward. In fact, two distinct approaches to this extension emerge, distinguished by the *timing* at which optimality is evaluated.

The first approach defines dominance as guaranteed optimality at the ex post stage. A strategy is dominant in this sense if it maximizes an agent's ex post (realized) utility against all possible strategies chosen by other agents. In other words, such a strategy remains optimal regardless of the strategies chosen by others, holding for every realization of their types. We refer to this notion as *ex post dominance*, and a mechanism in which every agent has an ex post dominant strategy is termed *ex post strategy-proof (ESP)*.

The second approach defines dominance as guaranteed optimality at the interim stage. A strategy is dominant in this sense if it maximizes an agent's interim (expected) utility against all possible strategies chosen by other agents. In other words, such a strategy remains optimal regardless of the strategies chosen by others, averaging over the distribution of their types. We refer to this notion as *interim dominance*, and a mechanism in which every agent has an interim dominant strategy is termed *interim strategy-proof (ISP)*.

In private value settings, both ESP and ISP coincide with the conventional notion of strategy-proofness. In the presence of preference interdependence, however, ESP is strictly stronger than ISP. This paper focuses on mechanism design under the more demanding ESP criterion, while a companion paper (Feng and Wu (2025)) investigates ISP mechanism design.

A central appeal of ESP mechanisms lies in their simplicity and robustness from the perspectives of both agents and designers. For agents, ESP mechanisms are user-friendly because optimal behavior does not require forming correct conjectures about others' strategies or the distribution of their types. In environments with interdependent values, such conjectures are typically indispensable for strategic optimization, as payoffs depend directly on these factors. In contrast, an ESP mechanism equips each agent with a strategy that is optimal against all possible strategies and type profiles of opponents. Consequently, even agents who are uninformed or lack strategic sophistication can reliably identify and play an optimal strategy. This reduces cognitive demands, limits the scope for strategic errors, and ensures that agents who lack the information or sophistication required to game the mechanism are not disadvantaged relative to more sophisticated one.

From the designer's perspective, ESP mechanisms enhance predictability, as dominant strategies are likely to be played even when agents lack the experience to learn about

others' strategies or the ability to coordinate on equilibrium behavior. In this sense, ESP mechanisms offer robustness to strategic uncertainty—that is, uncertainty regarding the beliefs agents hold about others' strategies. Such strategic robustness is particularly valuable in environments with many inexperienced participants unlikely to hold accurate beliefs, such as presidential elections involving a general electorate or large-scale school-choice procedures (Chen and Kesten, 2017). Moreover, because ESP mechanisms admit ex post equilibria, they are informationally robust in the sense of Bergemann and Morris (2005): their performance does not rely on the designer's assumptions about agents' beliefs or higher-order beliefs. This informational robustness is crucial when the designer faces substantial uncertainty about the underlying information structure.

However, these desirable properties of ESP mechanisms may come at a cost: the demanding nature of ESP may restrict the set 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 in environments with transfers.

We begin with a specific class of environments satisfying a property we term "one-dimensionality". This property extends an analogous concept from private value settings, where both alternatives and agent types can be ranked such that a higher-ranked type exhibits a stronger preference for a higher-ranked alternative than 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—requiring higher-ranked types to receive (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 term "monotonicity", remains necessary for ESP implementation.

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

The severity 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, screening different types becomes difficult. If types form a continuum, even minimal preference interdependence renders them mutually unscreenable, thereby severely limiting implementation. Conversely, when types are relatively sparse, entanglement may not occur despite the presence of preference interdependence, allowing any social choice function that satisfies monotonicity to be implementable.

Within the class of single-good auctions, we also study the implementability of efficient allocations. We show that efficiency is ESP-implementable if and only if a strengthened "strong single-crossing" condition holds, reflecting the additional constraints imposed

by preference interdependence.

We then examine the most general setting by relaxing the assumption of one-dimensionality. We establish that a social choice function is implementable if and only if it satisfies a condition generalizing cyclical monotonicity. While standard cyclical monotonicity characterizes strategy-proof implementation in private value settings (Rochet (1987)), a comparison with our generalized version highlights the additional incentive challenges introduced by preference interdependence. Specifically, the mechanism must ensure that misreporting is unprofitable for an agent, even when she evaluates the outcome using the most optimistic beliefs about other agents' types.

Additionally, we find that implementability can be characterized by an optimality condition in a two-sided matching problem between an agent's types and the alternatives. This result extends an analogous equivalence between strategy-proof implementation and optimal matching found in private value settings (Dworczak and Zhang (2017)).

Literature

The literature on strategy-proof mechanism design is extensive, yet it predominantly focuses on private value settings. While early discussions on strategy-proofness in interdependent value settings exist—notably in Crémer and McLean (1985) and Williams and Radner (1988)—a comprehensive study of this topic remains lacking. Instead, much of the literature concentrates on *ex post* implementation. A social choice function is ex post implementable if, in the corresponding direct mechanism, truthful reporting is optimal for every realization of the other agents' types, provided that those agents also report truthfully. Research suggests that nontrivial ex post implementation is achievable under specific 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 report untruthfully. Building on the ex post implementation literature, we find that conditions conducive to ex post implementation—such as one-dimensionality and single-crossing—can be suitably extended to support ESP implementation. However, ESP implementation is far more sensitive to preference interdependence. For instance, while many significant social choice functions are ex post implementable in one-dimensional environments even with a continuum type space,[3] they may fail to be ESP implementable if the type space lacks sufficient discreteness.

In a companion paper (Feng and Wu, 2025), we analyze *interim* strategy-proofness (ISP), a weaker extension of strategy-proofness tailored to interdependent value settings. ISP relaxes the requirement of ex post optimality, and this relaxation is substan-

[1] For example, Dasgupta and Maskin (2000) require one-dimensional types and a single-crossing condition for efficient implementation, while Bikhchandani (2006) allow for multi-dimensional types but require 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).

tive. Generally, ISP allows for the implementation of a broader range of social choice functions than ESP. For instance, in binary voting, many nontrivial social choice functions are ISP implementable regardless of the preference or information structure;in contrast, as we show in this paper, this is not the case under ESP.

Our analysis of the general setting extends the findings of Rochet (1987) to interdependent value settings. We identify a generalization of Rochet's cyclical monotonicity condition that provides an exact characterization of ESP implementability, mirroring how standard cyclical monotonicity characterizes strategy-proofness in private value settings. Additionally, we derive a generalization of the one-dimensionality condition for interdependent values. We demonstrate that, under this condition, generalized cyclical monotonicity is equivalent to monotonicity in the social choice function modulo entangled types. This finding yields the analogous result in Myerson (1981) (modified for strategy-proof implementation) as a special case.

ESP implementation shares conceptual similarities with the problem of maxmin implementation studied in Tang and Zhang (2021). In both contexts, agents respond conservatively to uncertainty. In our framework, uncertainty relates to the strategies and types of other agents, whereas in Tang and Zhang (2021), it arises from ambiguity regarding the mechanism's execution. Both papers characterize implementation using a version of cyclical monotonicity, though the specific formulations and interpretations differ.

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

Regarding our analytical approach, while we work directly with incentive constraints, the problem admits an equivalent network-flow formulation—in the spirit of Vohra (2011) and Rahman (2024)—where arc lengths between types are defined by incentives. Under this formulation, some of our results can be derived analogously.[4]

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 alternative x to agent i given the type profile θ, and t_i is the transfer she pays to the mechanism. The function u_i may depend on other agents' types θ_{-i}, so the environment allows for interdependent values. Agents evaluate lotteries over outcomes according to expected utility.

[4]We thank an anonymous referee for this observation.

A mechanism (with transfers) consists of message sets $M_1, \ldots, M_n$, one for each agent; an outcome function $q : M_1 \times \ldots \times M_n \to X$; and a transfer function $\tau = (\tau_1, \ldots, \tau_n) : M_1 \times \ldots \times M_n \to \mathbb{R}^n$. In this mechanism, each agent i reports a message $m_i \in M_i$. The resulting report profile $m = (m_1, \ldots, m_n)$ determines the chosen alternative $q(m)$ and the payment $\tau_i(m)$ required from each agent i.

We say a strategy $\sigma_i : \Theta_i \to M_i$ for agent i is *ex post dominant* if, for every type $\theta_i \in \Theta_i$, the utility from $\sigma_i(\theta_i)$ is weakly higher than the utility from any alternative $\sigma'_i(\theta_i)$, regardless of the other agents' realized types and strategies.

Our objective is to identify which social choice functions—mappings from Θ to X—are implementable in an equilibrium where every agent plays an ex post dominant strategy. Invoking the revelation principle,[5] we restrict our attention to *direct mechanisms*, in which $M_i = \Theta_i$ for every agent i.

Definition 2.1. A direct mechanism with outcome function q and transfer function τ is *ex post strategy-proof (ESP)* if, for every agent i, every $\theta_i, \theta'_i \in \Theta_i$, and every $\theta_{-i}, \theta'_{-i} \in \Theta_{-i}$,

$$u_i\big(q(\theta_i, \theta'_{-i});\, \theta_i, \theta_{-i}\big) - \tau_i(\theta_i, \theta'_{-i}) \;\geq\; u_i\big(q(\theta'_i, \theta'_{-i});\, \theta_i, \theta_{-i}\big) - \tau_i(\theta'_i, \theta'_{-i}).$$

Thus, a direct mechanism is ESP if truthful reporting of types constitutes an ex post dominant strategy for every agent. If u_i does not depend on θ_{-i}, the environment is one of private values, and ESP reduces to standard strategy-proofness. In this case, strategy-proofness is known to coincide with ex post incentive compatibility (EPIC). By contrast, in environments with interdependent values, ESP is generally stronger than EPIC: while ESP requires that the above inequality hold for all reported type profiles θ'_{-i} of the other agents, EPIC requires it only when the others report truthfully, i.e. when $\theta'_{-i} = \theta_{-i}$.

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 (q, τ) is ESP.

3 Analysis

We begin our analysis with a simple auction example, which provides a preview of a more general single-good auction model. We then broaden our scope to a class of "one-dimensional" environments—a category that encompasses auctions as a special case. Finally, we characterize implementation within the most general setting. All proofs are provided in the Appendix.

3.1 Leading example

Consider a seller looking to sell a single indivisible good to two interested buyers, agents $i = 1, 2$. Each agent i has a type represented by a real number θ_i, which takes values from the set $\Theta_i = \{0, \frac{1}{K}, \frac{2}{K}, \cdots, \frac{K}{K}\}$, where K is a positive natural number. Observe that the types are evenly distributed between 0 and 1, with K serving as a parameter for type density. As K increases, the gap between neighboring types $(1/K)$ decreases, resulting

[5] The proof follows standard arguments and is therefore omitted.

in a denser set of types. Consequently, 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 reflecting the degree of preference interdependence. A higher value of β indicates that the other agent's type exerts a greater influence on an agent's valuation. Specifically, when $\beta = 0$, the model represents a pure private value auction; conversely, when $\beta = 1$, it represents a pure common value auction.

If, as assumed in many studies, the joint distribution of the agents' types is common knowledge among the agents and the designer, and if agents can be relied upon to correctly conjecture each other's bidding behavior, then the designer can employ a Bayesian mechanism to achieve his objective. However, if the designer cannot safely assume that agents know each other's bidding strategies—perhaps because they lack interaction experience—then an interim strategy-proof mechanism is suitable, as it eliminates the need for such strategic knowledge. Alternatively, if the designer cannot assume that agents know the type distribution, or if the seller himself is ignorant of this distribution, then an ex post mechanism is appropriate. Finally, an ESP mechanism is suitable when the designer cannot assume that agents possess accurate knowledge of either one another's strategies or the type distribution.

A direct mechanism is defined 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, $q_i(\theta)$ represents the probability that agent i wins the good given the reported type profile θ, and $\tau_i(\theta)$ denotes her corresponding payment.

Which allocation functions are implementable? We show that the answer depends on the interplay between the parameters K and β. Specifically, if $\beta K > 1$, then q is implementable if and only if each q_i is constant with respect to 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, consider an agent i and two of her neighboring types, θ_i and $\theta_i' = \theta_i - 1/K$. Incentive compatibility requires that agent i has no incentive to report type θ_i' when her true type is θ_i, regardless of the other agent's actual type θ_j and reported type θ_j'. Formally, for all $\theta_j, \theta_j' \in \Theta_j$:

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

Similarly, ensuring that type θ_i' has no incentive to report type θ_i requires that for all $\theta_j, \theta_j' \in \Theta_j$:

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

The standard argument—adding the two inequalities and rearranging—yields $q_i(\theta_i, \theta_j') \geq q_i(\theta_i', \theta_j')$. This establishes the familiar monotonicity condition: the winning probability must be weakly increasing in the agent's type.

Can $q_i(\theta_i, \theta_j')$ be *strictly* greater than $q_i(\theta_i', \theta_j')$? To answer this, we analyze inequalities (1) and (2) considering all possible values of $\theta_j \in \Theta_j$. From inequality (1), since the

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

condition must hold for all θ_j, we have for all $\theta'_j \in \Theta_j$:

$$\big(q_i(\theta_i, \theta'_j) - q_i(\theta'_i, \theta'_j)\big) \min_{\theta_j \in \Theta_j} v_i(\theta_i, \theta_j) \geq \tau_i(\theta_i, \theta'_j) - \tau_i(\theta'_i, \theta'_j). \tag{3}$$

Similarly, from inequality (2), we have for all $\theta'_j \in \Theta_j$:

$$\tau_i(\theta_i, \theta'_j) - \tau_i(\theta'_i, \theta'_j) \geq \big(q_i(\theta_i, \theta'_j) - q_i(\theta'_i, \theta'_j)\big) \max_{\theta_j \in \Theta_j} v_i(\theta'_i, \theta_j). \tag{4}$$

Combining these bounds, it is clear that $q_i(\theta_i, \theta'_j)$ can be strictly greater than $q_i(\theta'_i, \theta'_j)$ for some θ'_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 $q_i(\theta_i, \theta'_j)$ must equal $q_i(\theta'_i, \theta'_j)$ for all θ'_j. In other words, θ_i and θ'_i cannot result in 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 does not guarantee that type θ_i can achieve a strictly higher winning probability in an ESP mechanism. Instead, a more stringent condition is necessary: the lowest possible value under θ_i must be no less than the highest possible value under θ'_i.[9] Otherwise, θ_i and θ'_i must result in the same winning probability.

Screenability plays a crucial role in ESP-implementation. To better understand when types are screenable, let us consider the *valuation interval* for each type θ_i, defined as:

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

The valuation interval represents the smallest closed interval in $\mathbb{R}$ containing 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 knows her value must fall within this interval, regardless of the realization of θ_j.

Now, returning to inequality (5), a necessary condition for screenability, we can visualize its satisfaction and violation as follows:

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

[8] Note that even if the two types are unscreenable from agent i's perspective, they may still imply different winning probabilities for agent j. Such "non-self-responsive" mechanisms are of general interest. For instance, Niemeyer and Preusser (2024) show that when agents are "informationally small," a class of non-self-responsive mechanisms—termed "ranking-based" mechanisms—achieve approximate optimality in certain environments without transfers.

[9] When this condition is satisfied, agent i's preferences remain genuinely interdependent: her valuation for the good—and hence her willingness to pay— varies with θ_j. The condition requires only that the lowest possible willingness to pay under type θ_i weakly exceed the highest possible willingness to pay under θ'_i.

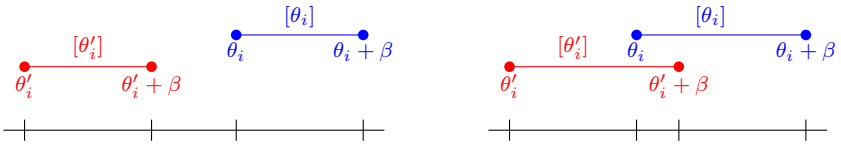

Inequality (5) satisfied

Inequality (5) violated

We observe that inequality (5) is violated (rendering the two types unscreenable) if $[\theta_i]$ and $[\theta_i']$ have a non-degenerate intersection—that is, an intersection that is both nonempty and non-singleton. In this case, we say that θ_i and θ_i' *overlap*. Numerically, this occurs when $\beta K > 1$. Graphically, this happens when the two valuation intervals widen enough to bridge the gap between them and intersect. This indicates that when preference interdependence (measured by β) is high relative to the density of types (measured by K), neighboring types begin to overlap and become unscreenable.

Inequality (5) is satisfied when $[\theta_i]$ lies entirely to the right of $[\theta_i']$ with no non-degenerate intersection. In this case, θ_i *might* be screened from θ_i'—meaning there could exist an implementable allocation function q such that $q_i(\theta_i, \theta_j') > q_i(\theta_i', \theta_j')$ for some θ_j'. However, unscreenability remains possible. To illustrate, suppose $\beta K > 1$, implying that neighboring types overlap. Consider two non-neighboring types θ_i and θ_i' with $\theta_i > \theta_i'$ that do not themselves 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 on, eventually reaching θ_i'. In other words, we can "travel" from θ_i to θ_i' via a finite sequence of types, where each overlaps with—and is therefore unscreenable from—the next. In such instances, we describe types θ_i and θ_i' as *entangled*. Since unscreenability is transitive, two entangled types are unscreenable even if they do not directly overlap. 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':

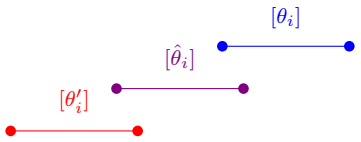

Conversely, if two types are not entangled, they are screenable. In the current example, this occurs when $\beta K \leq 1$, as even neighboring types do not overlap. In this scenario, any allocation function q such that q_i is non-decreasing in θ_i for each i is implementable. The supporting transfer function follows a standard construction: for any θ_j', we set $\tau_i(0, \theta_j') = 0$, and for each $\theta_i = \frac{1}{K}, \ldots, \frac{K}{K}$, we recursively define:

$$\tau_i(\theta_i, \theta_j') = \tau_i(\theta_i - 1/K, \theta_j') + \theta_i\big(q_i(\theta_i, \theta_j') - q_i(\theta_i - 1/K, \theta_j')\big).$$

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

- Monotonicity: If two types of the same agent are not entangled, they are screenable, allowing the higher type to achieve a strictly higher winning probability.
- Screenability: If two types of the same agent are entangled (including the special case where they overlap), they are not screenable.

3.2 Single-good auctions

We now analyze a general single-good auction model. There are n buyers (agents) indexed by $i = 1, \ldots, n$, 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 varies with the types of other agents. Formally, this entails 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 $v_i(\theta_i, \theta_{-i}) \neq v_i(\theta_i, \theta'_{-i})$. [10]

A direct mechanism is given by an allocation function $q = (q_1, \ldots, q_n) : \Theta \to \Delta^n$ [11] and a transfer function $\tau = (\tau_1, \ldots, \tau_n) : \Theta \to \mathbb{R}^n$.

Implementability

We generalize the leading example by establishing that monotonicity and screenability jointly characterize implementability in the single-good auction model. The *valuation interval* of a type $\theta_i \in \Theta_i$ is

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

Definition 3.1. For any agent i and any $\theta_i, \theta'_i \in \Theta_i$:

- θ_i and θ'_i *overlap* if their valuation intervals have a non-degenerate intersection.
- θ_i and θ'_i are *entangled* if there exists a finite sequence of types in Θ_i linking θ_i to θ'_i such that consecutive types in the sequence overlap.
- θ_i is *higher than* θ'_i if (1) θ_i and θ'_i are not entangled, and (2) the valuation interval $[\theta_i]$ lies weakly to the right of $[\theta'_i]$, that is, $\max[\theta'_i] \leq \min[\theta_i]$.

As in the leading example, two types overlap when their respective valuations cannot be unambiguously ranked without knowing the types of other agents. Furthermore, types are entangled if they overlap directly or are bridged by a chain of overlapping intermediate types.

Proposition 3.1. *An allocation function q in the single-good auction is implementable if and only if for any agent i, any $\theta_i, \theta'_i \in \Theta_i$, and any $\theta_{-i} \in \Theta_{-i}$, the following conditions hold:*

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

This proposition follows as a corollary of a more general result presented in the next section. The heuristic analysis provided in the leading example offers a rough outline of the proof.

[10] 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.

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

One important implication of Proposition 3.1 is that, as illustrated in the leading example, the scope of ESP implementation is jointly constrained by two factors: the *density of types* (how closely valuation intervals are situated to one another) and the *extent of preference interdependence* (the widths of those valuation intervals). When types are densely distributed or when preference interdependence is significant, valuation intervals begin to overlap. This renders types unscreenable, thereby limiting the ability of an allocation function to differentiate between them.

If Θ_i is a connected subset of $\mathbb{R}^d$ and v_i is continuous in θ_i, it follows that as long as there is "minimal" preference interdependence—in the sense that the widths of valuation intervals are uniformly bounded away from zero—no two types can be screened.[12] 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 does not hinder implementation, provided it is not so significant as to cause valuation intervals to overlap. This is exemplified in the leading example when $\beta K \leq 1$; in this case, 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 a model in which types are drawn from $\Theta_i = [0,1]^{d_i}$, and establish a generic impossibility for the existence of non-trivial EPIC mechanisms. In light of our findings, this impossibility can be understood as hinging on the fact that the type space forms a continuum, so that types are densely distributed. In contrast, when types are not as densely distributed—for example, when they are discrete as in our leading example—nontrivial ESP mechanisms, which are *a fortiori* also EPIC, can exist even with multi-dimensional types.

Efficiency

Next, we investigate the implementability of efficient allocation functions—those that allocate the good to the agent who values it most. For simplicity, we assume that no two agents assign the same value to the good: for any $\theta \in \Theta$, $v_i(\theta) \neq v_j(\theta)$ whenever $i \neq j$. Let $i^*(\theta)$ denote the unique agent with highest valuation at type profile θ. There is a unique efficient allocation function q^E:

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

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, single-crossing stipulates that if an agent has the highest valuation, she maintains that status if her type increases (assuming real-valued types) while others' types remain fixed.

We show that a stronger, yet conceptually related, condition is necessary and sufficient

[12] Since we maintain the assumption that Θ_i is finite for every agent i, a more rigorous statement would consider a sequence of environments with finite type spaces that become increasingly dense and approximate a connected subset of $\mathbb{R}^d$. The leading example illustrates this idea by letting $K \to \infty$ while β remains bounded away from zero.

for the ESP implementation of q^E. We define θ_i' as being *weakly higher* than θ_i if the two types are entangled or if θ_i' is strictly higher than θ_i.

Definition 3.2. A single-good auction environment satisfies *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 directly compare single-crossing with strong single-crossing, we adopt the setting of Dasgupta and Maskin (2000), assuming types are real numbers and 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' greater than θ_i, but also for θ_i' smaller than θ_i as along as θ_i and θ_i' are entangled.

Strong single-crossing implies that for any θ_{-i}, the types for which agent i is the highest valuer must be screenable from those for which she is not. In our leading example, when $\beta K > 1$, all types of a given agent are entangled. Consequently, strong single-crossing would require that an agent's status as the highest valuer be independent of her type—a condition that is clearly violated. Indeed, in that example, strong single-crossing holds if $\beta \leq 1/K$, whereas single-crossing holds 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.*

The necessity of strong single-crossing is straightforward. By Proposition 3.1, if agent i wins under reported type profile θ, she must continue to win if her type becomes weakly higher. Since q^E allocates the good to the highest valuer, implementation requires that a highest-valuing agent maintains that status across all weakly higher types. This is precisely the definition of strong single-crossing.

If strong single-crossing holds, we show in the proof that q^E is implemented by transfers equivalent to the following pricing scheme:

- If agent i does not win, she pays a baseline fee $\underline{p}_i(\theta_{-i})$.
- If agent i wins, she pays the baseline fee $\underline{p}_i(\theta_{-i})$ plus an additional amount $p_i(\theta_{-i})$ satisfying:

$$\max_{\tilde{\theta} \in \Theta:\ i \neq i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}) \ \leq \ p_i(\theta_{-i}) \ \leq \ \min_{\tilde{\theta} \in \Theta:\ i = i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}). \tag{6}$$

In fact, any direct mechanism implementing q^E must imply precisely such a pricing scheme. Strong single-crossing guarantees the existence of a $p_i(\theta_{-i})$ satisfying the inequalities in (6). One specific example is:

$$\underline{p}_i(\theta_{-i}) = 0 \quad \text{and} \quad p_i(\theta_{-i}) = \max_{\tilde{\theta} \in \Theta:\ i \neq i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}).$$

Under this scheme, losing agents pay nothing. The winner pays an amount equal to the highest value she could have assigned to the good while still losing. This pricing structure extends the ex post implementation scheme of Dasgupta and Maskin (2000) and is conceptually analogous to the Vickrey auction.

Williams and Radner (1988) demonstrates that when types form a real interval, imple-

menting q^E via an ISP mechanism is impossible given even minimal preference interdependence. Our findings suggest this impossibility hinges on the fact that the type space forms a continuum, with no discreteness. As Proposition 3.2 and the leading example (where $\beta K \leq 1$) illustrate, when types are less densely distributed, mild preference interdependence does not preclude the ESP—and thus ISP—implementation of q^E.

3.3 Generalized one-dimensional environments

Single-good auctions represent a special case within a broader class of environments where implementability is jointly characterized by monotonicity and screenability. We call this class *generalized one-dimensional*, as it encompasses standard one-dimensional environments in the absence of preference interdependence.[13]

We begin by defining this class of environments. For each agent i, let

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

denote the marginal value of alternative x over alternative y, given type profile θ. Furthermore, let

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

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 regarding θ_{-i}. We refer to $[\theta_i]^{xy}$ as a *marginal value interval*.

The valuation intervals introduced for single-good auctions are a special case of the marginal value intervals defined below, which extend this notion to arbitrary pairs of alternatives. In single-good auctions, an alternative in the set X is represented by a vector of winning probabilities $p = (p_1, \ldots, p_n)$. The payoff to agent i from any $p \in X$ is given by $u_i(p;\theta) = p_i v_i(\theta)$. Consequently, 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 versus p' is therefore:

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

In other words, $[\theta_i]^{pp'}$ is obtained from the valuation interval

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

by multiplying its endpoints by $p_i - p_i'$, with the ordering of the bounds adjusted when $p_i - p_i' < 0$. A useful observation is that two types θ_i and θ_i' overlap if and only if, for any

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

pair $p, p' \in X$ with $p_i \neq p'_i$, the corresponding marginal value intervals $[\theta_i]^{pp'}$ and $[\theta'_i]^{pp'}$ have a non-degenerate intersection.

As in the leading example and single-good auctions, the concept of entanglement in generalized one-dimensional environments is defined based on the relative positions of intervals on the real line. Let A and B be two closed intervals. We define the following relations:

- $A \to B$ 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 $A \nrightarrow B$ and $B \nrightarrow A$.

The notation $A \to B$ indicates that interval B lies entirely to the right of interval A, with no non-degenerate intersection. Specifically, depending on whether the intervals are degenerate (i.e., singletons) and whether they share endpoints, one of the following scenarios holds:

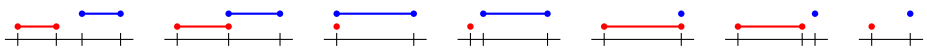

$A \to B$

The notation $A \bowtie B$ generalizes the overlapping relation between two intervals and captures the idea that the two intervals cannot be unambiguously ordered. This occurs in three cases:

1. If both A and B are non-degenerate, their intersection is non-degenerate.

2. If one interval is degenerate and the other is non-degenerate, the degenerate interval lies strictly in the interior of the non-degenerate one.

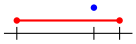

3. If both A and B are degenerate, they coincide.

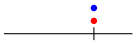

We now formalize this notion of entanglement for general environments.

Definition 3.3. Let $\sim^X$ be an equivalence relation on the set of alternatives X. For any agent i, $\theta_i, \theta'_i \in \Theta_i$ are said to be *generalized entangled with respect to* $\sim^X$ if there exists a finite sequence of types $(\theta_i^k)_{k=1}^K$ in Θ_i, with $\theta_i^1 = \theta_i$ and $\theta_i^K = \theta'_i$, such that for every $k \in \{1, \ldots, K-1\}$ and every pair of alternatives $x, y \in X$ satisfying $x \not\sim^X y$,

$$[\theta_i^k]^{xy} \bowtie [\theta_i^{k+1}]^{xy}.$$

In generalized one-dimensional environments, two types of agent i are considered generalized entangled if they can be connected by a sequence of intermediate types. This connection requires that for any pair of alternatives that are non-equivalent to agent i, the marginal value intervals associated with that pair must consecutively overlap

(having the relation $\bowtie$) along the sequence. This general definition encompasses the concept of entanglement in single-good auctions. For example, if we define the equivalence relation $\sim^X$ such that $p \sim^X p'$ whenever $p_i = p'_i$, the general definition implies the specific case: if the marginal value intervals for a specific pair p, p' overlap for two types, it follows that the intervals for *any* pair p, p' with $p_i \neq p'_i$ must also overlap for those types.

Definition 3.4. An environment satisfies *generalized one-dimensionality (G1D)* if, for each agent i, there exist complete and transitive relations $\succsim_i^X$ on X and $\succsim_i^\Theta$ on Θ_i (with $\succ_i^X$ / $\succ_i^\Theta$ and $\sim_i^X$ / $\sim_i^\Theta$ denoting the respective induced strict and equivalence relations) such that the following conditions hold:

(G1) *Increasing differences.* If $x \succ_i^X y$ and $\theta_i \succ_i^\Theta \theta'_i$, then $[\theta'_i]^{xy} \to [\theta_i]^{xy}$.

(G2) *Indifference between tied alternatives.* If $x \sim_i^X y$, then $[\theta_i]^{xy} = \{0\}$ for all $\theta_i \in \Theta_i$.

(G3) *Tied types.* If $\theta_i \sim_i^\Theta \theta'_i$ and $\theta_i \neq \theta'_i$, then:

- (a) *Entanglement.* θ_i and θ'_i are generalized entangled with respect to $\sim_i^X$.
- (b) *No redundancy.* For every pair $x \not\sim_i^X y$, the set $[\theta_i]^{xy} \cup [\theta'_i]^{xy}$ cannot be a singleton.

To interpret G1D, consider standard one-dimensionality in private value settings. There, one-dimensionality implies that for each agent, both alternatives and types can be ordered from "low" to "high" such that the valuation function exhibits increasing differences: higher types have (weakly) higher marginal values for higher-ranked alternatives. Because marginal values are independent of θ_{-i}, types can be unambiguously ordered.

With interdependent values, however, marginal values generally depend on θ_{-i}, rendering the ordering of types potentially ambiguous: θ_i may have a higher marginal value for x over y than θ'_i for some θ_{-i}, yet a lower value for others. Condition (G1) addresses this by strengthening the concept of increasing differences to require an *unambiguous* ordering of types: when comparing a higher alternative x versus a lower alternative y, for θ_i to be a higher type than θ'_i, the worst-case marginal value under θ_i must dominate the best-case marginal value under θ'_i (i.e., $[\theta'_i]^{xy} \to [\theta_i]^{xy}$). Condition (G2) preserves the standard requirement that two alternatives are tied under $\succsim_i^X$—and should be considered payoff-equivalent—if every type is indifferent between them.

Condition (G3) is specific to interdependent values and highlights how ties in the type order arise due to preference interdependence. Specifically, types can be tied only when preference interdependence entangles them in comparisons of alternatives that are not payoff-equivalent. Condition (G3.a) formalizes this by requiring that tied types belong to the same entanglement class—meaning they are linked by a finite chain of types with overlapping marginal value intervals for *every* binary comparison $x \not\sim_i^X y$. Entanglement is, therefore, the obstruction to strictly ordering types. G1D preserves a one-dimensional structure of types by permitting ties only within these entanglement classes; across classes, types are strictly ordered consistently across all comparisons. This structure is crucial for implementation, as we formally show shortly: entangled types cannot be screened due to incentive constraints and must be treated identically (implying the *screenability* condition), while strictly ordered types can be screened

monotonically (implying the *monotonicity* condition).

Finally, Condition (G3.b) is a mild restriction imposed for expositional simplicity. At an intuitive level, it rules out the possibility that two tied types are *redundant private-value types* with respect to any nontrivial comparison (i.e., $x \not\sim_i^X y$). Specifically, if the union of marginal value intervals $[\theta_i]^{xy} \cup [\theta_i']^{xy}$ were a singleton, then both intervals would necessarily coincide as that same singleton. This would imply that these two types behave as if values were private and identical for that specific comparison. While allowing such redundancies would not alter the fundamental analysis, it would complicate the statement of results (e.g., Proposition 3.3) by permitting multiple type orders to satisfy G1D. In private value settings, types cannot be entangled because all marginal value intervals are degenerate. Consequently, once redundant types are excluded by Condition (G3.b), standard one-dimensionality is equivalent to G1D.

To gain a more concrete understanding of what G1D entails, it is useful to consider the following example.

Example 3.1. Let $X \subset \mathbb{R}$ be a convex set, and let $\Theta_i \subset \mathbb{R}$ for each agent $i = 1, \dots, n$. For each agent i, u_i is continuously differentiable in x with a strictly positive partial derivative on $\operatorname{int}(X)$. In addition, u_i satisfies the following conditions for any $\theta_i^h, \theta_i^l \in \Theta_i$ with $\theta_i^h > \theta_i^l$:

1. *Spence–Mirrlees.*

$$\forall \theta_{-i} \in \Theta_{-i} : \quad \frac{\partial}{\partial x} u_i(x; \theta_i^h, \theta_{-i}) > \frac{\partial}{\partial x} u_i(x; \theta_i^l, \theta_{-i}) \qquad \forall x \in \operatorname{int}(X).$$

2. *Additional requirement for G1D.* One of the following must hold:

$$\text{(2a)} \quad \forall\, \theta_{-i}, \theta_{-i}' \in \Theta_{-i} : \quad \frac{\partial}{\partial x} u_i(x; \theta_i^h, \theta_{-i}) \geq \frac{\partial}{\partial x} u_i(x; \theta_i^l, \theta_{-i}') \qquad \forall x \in \operatorname{int}(X),$$

$$\text{(2b)} \quad \exists\, \theta_{-i}, \theta_{-i}' \in \Theta_{-i} : \quad \frac{\partial}{\partial x} u_i(x; \theta_i^h, \theta_{-i}) < \frac{\partial}{\partial x} u_i(x; \theta_i^l, \theta_{-i}') \qquad \forall x \in \operatorname{int}(X).$$

In Appendix A.3, we verify that this example satisfies G1D.

Condition (1) is the standard Spence–Mirrlees single-crossing property: agent i's marginal value under a high type θ_i^h always exceeds that under a low type θ_i^l, fixing any type profile θ_{-i} of the others.

Condition (2) imposes a dichotomy on the separation of types. It requires that for any pair of types, either:

(2a) the high type is sufficiently "separated" from the low type, such that the high-type marginal value is unambiguously higher than the low-type marginal value, even if they are evaluated at different type profiles of others; or

(2b) the high type is sufficiently "close" to the low type, such that the high-type marginal value is *never* unambiguously higher than the low-type marginal value when evaluated at possibly different type profiles of others.

Single-good auctions provide another example of G1D environments. Specifically, the conditions in the definition of G1D hold[14] for the following orders of alternatives:

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

and types:

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

Consequently, the relation $\theta_i \sim_i^\Theta \theta_i'$ holds if θ_i and θ_i' are entangled, whereas $\theta_i \succ_i^\Theta \theta_i'$ holds if θ_i is strictly higher than θ_i'. (These relations—entanglement and being higher than—are formally defined in Section 3.2.)

Therefore, Proposition 3.1 can be reformulated to state that in single-good auctions, a social choice function is implementable if and only if a higher type yields a weakly higher alternative, while entangled types yield tied alternatives. This principle generalizes to any environment satisfying G1D:

Proposition 3.3. *Suppose an environment satisfies generalized one-dimensionality with orders $(\succsim_i^X, \succsim_i^\Theta)_{i=1:n}$ over alternatives and types. A social choice function q is implementable if and only if for every agent i, any $\theta_i, \theta_i' \in \Theta_i$, and any $\theta_{-i} \in \Theta_{-i}$, the following conditions hold:*

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

3.4 Implementability in general

In this section, we analyze the general environment outlined in Section 2. Specifically, we present two characterizations of implementable social choice functions.

First characterization: generalized cyclical monotonicity

The first characterization is formulated in terms of an extension of the well-known cyclical monotonicity condition[15] to interdependent value settings. This extension is defined as follows:

Definition 3.5. A social choice function q satisfies *generalized cyclical monotonicity (GCM)* if for every agent i, every $\theta_{-i}' \in \Theta_{-i}$, and every finite sequence of agent i's types $(\theta_i^k)_{k=1}^K$ such that $\theta_i^1 = \theta_i^K$, we have

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

To interpret this, observe that the term in the braces represents the utility gain (modulo transfers) for agent i of type θ_i^k who deviates by reporting θ_i^{k+1} instead of her true type. GCM is satisfied if, for any closed chain of such deviations, the sum of these "most favorable" utility gains is non-positive.

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

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

Note that if u_i does not vary with θ_{-i}—as is the case with private values—the $\max$ operator becomes redundant. In that case, GCM reduces 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 of 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 must consider all possible realized type profiles of the other agents. Incentive compatibility requires that a deviation not be profitable even under the type profile of others that is most favorable to the deviation—hence the inclusion of the $\max$ operator. Crucially, this $\max$ operator is applied term-by-term within the summation over the deviation chain. This reflects the fact that a cycle is composed of distinct deviations, each potentially evaluated against a different "most favorable" profile θ_{-i}.

We note that for single-value auctions—and more generally for G1D environments—the necessity of monotonicity and screenability for implementation is immediately implied by GCM by considering deviation chains of length two. To see this, consider a single-good auction, i.e. $u_i(q;\theta) = q_i v_i(\theta)$. Fix an agent i and two types θ_i, θ_i', and consider the deviation chain $\theta_i \to \theta_i' \to \theta_i$. GCM requires that for any $\theta'_{-i} \in \Theta_{-i}$:

$$\max_{\theta_{-i}} \left\{ \Big(q_i(\theta_i', \theta'_{-i}) - q_i(\theta_i, \theta'_{-i}) \Big) v_i(\theta_i, \theta_{-i}) \right\} + \max_{\theta_{-i}} \left\{ \Big(q_i(\theta_i, \theta'_{-i}) - q_i(\theta_i', \theta'_{-i}) \Big) v_i(\theta_i', \theta_{-i}) \right\} \leq 0.$$

Suppose $q_i(\theta_i, \theta'_{-i}) > q_i(\theta_i', \theta'_{-i})$. The inequality rearranges as:

$$-\Big(q_i(\theta_i, \theta'_{-i}) - q_i(\theta_i', \theta'_{-i}) \Big) \min_{\theta_{-i}} v_i(\theta_i, \theta_{-i}) + \Big(q_i(\theta_i, \theta'_{-i}) - q_i(\theta_i', \theta'_{-i}) \Big) \max_{\theta_{-i}} v_i(\theta_i', \theta_{-i}) \leq 0,$$

which implies

$$\max_{\theta_{-i}} v_i(\theta_i', \theta_{-i}) \leq \min_{\theta_{-i}} v_i(\theta_i, \theta_{-i}).$$

This implies that a strictly higher winning probability can only be assigned to a higher type, necessitating monotonicity. It also follows that overlapping types cannot be assigned distinct winning probabilities. Given the transitivity of the overlap relation, entangled types must therefore be treated identically, necessitating screenability.

Second characterization: optimal matching

The second characterization differs significantly in nature. It establishes a connection between social choice functions and matching problems, demonstrating 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, the implementability of a social choice function is equivalent to the utilitarian efficiency of an induced matching between the agent's types and the alternatives, treating each type as a distinct individual. Our characterization extends their findings to settings with multiple agents and

interdependent values.

To formalize this, fix a social choice function q, an agent i, and a report profile θ'_{-i} of the other agents. We can view the function $q(\cdot, \theta'_{-i})$ as defining a matching $\mu : \Theta_i \to X$ between the types of agent i and the alternatives in X. Specifically, each type θ_i of agent i is matched to an alternative given by $\mu(\theta_i) := q(\theta_i, \theta'_{-i})$. 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 a *type-alternative matching for agent* i.

Given any type-alternative matching $\mu : \Theta_i \to X$ for agent i, another type-alternative matching $\mu' : \Theta_i \to X$ 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 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 $D_i(\mu', \mu)$, observe that the term in the curly brackets represents, for a given report 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 comparison is made under the scenario most favorable to μ', 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 $D_i(\mu', \mu) \leq 0$, 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 3.6. 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, we say that a social choice function q is reshuffling-efficient if for any agent i and any report profile $\theta'_{-i} \in \Theta_{-i}$, the induced type-alternative matching for agent i, $q(\cdot, \theta'_{-i})$, is reshuffling-efficient.

Reshuffling-efficiency embodies a stringent notion of utilitarian optimality. 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 stems 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.

Observe 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 the utilitarian efficiency 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 viewed 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 simply by 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 utilitarian welfare among all reshufflings of μ. Furthermore, this implies that a nontrivial implementable social choice function always exists. However, neither conclusion necessarily holds in interdependent value settings. To see this, consider the following example: For each agent i, each type θ_i, and any ordered pair of distinct alternatives x, 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 distinct from 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.

A Appendix

A.1 Proof of Proposition 3.1

(This proof relies on the analysis in Section 3.3.)

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

We show that the single-good auction model satisfies G1D with orders $(\succsim_i^X, \succsim_i^\Theta)_{i=1,\ldots,n}$ over alternatives and types. The completeness and transitivity of $\succsim_i^X$ are evident. The completeness of $\succsim_i^\Theta$ follows 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 partition of Θ_i induced by $\sim_i^\Theta$; that is, two types θ_i, θ'_i belong to the same element $P \in \mathcal{P}$ if and only if $\theta_i \sim_i^\Theta \theta'_i$. For each $P \in \mathcal{P}$, let $[P] := \bigcup_{\theta_i \in P} [\theta_i]$. Observe

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

that for any distinct $P, P' \in \mathcal{P}$, we have $[P] \not\bowtie [P']$. If this were not the case, there would exist $\theta_i \in P$ and $\theta_i' \in P'$ such that $[\theta_i] \bowtie [\theta_i']$, or equivalently $\theta_i \sim_i^\Theta \theta_i'$, which contradicts the assumption that $P \neq P'$.

Consider arbitrary types $\theta_i, \theta_i', \theta_i''$ such that $\theta_i \succsim_i^\Theta \theta_i'$ and $\theta_i' \succsim_i^\Theta \theta_i''$. Let $P, P', P'' \in \mathcal{P}$ be the equivalence classes containing these types, respectively. The condition $\theta_i \succsim_i^\Theta \theta_i'$ implies that either $P = P'$ or $[\theta_i'] \to [\theta_i]$ (which further implies $[P'] \to [P]$). Similarly, $\theta_i' \succsim_i^\Theta \theta_i''$ implies that either $P' = P''$ or $[P''] \to [P']$. We analyze the four possible cases:

1. If $P = P' = P''$, then $\theta_i \sim_i^\Theta \theta_i''$, which implies $\theta_i \succsim_i^\Theta \theta_i''$.
2. If $P = P'$ and $[P''] \to [P']$, then $[P''] \to [P]$. This implies $[\theta_i''] \to [\theta_i]$, and thus $\theta_i \succ_i^\Theta \theta_i''$, satisfying $\theta_i \succsim_i^\Theta \theta_i''$.
3. Similarly, if $[P'] \to [P]$ and $P' = P''$, then $[P''] \to [P]$, leading to $\theta_i \succsim_i^\Theta \theta_i''$.
4. If $[P'] \to [P]$ and $[P''] \to [P']$, then it follows that $[P''] \to [P]$, leading to $\theta_i \succsim_i^\Theta \theta_i''$.

We conclude that $\succsim_i^\Theta$ is transitive.

Next, we show that the environment manifests G1D. Note that for any θ_i and $p, p' \in X$,

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

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

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

We now verify the three conditions of G1D:

Condition (G1): 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 $[\theta_i']^{pp'} \to [\theta_i]^{pp'}$.

Condition (G2): Suppose $p \sim_i^X p'$. Then $p_i = p_i'$, and thus $[\theta_i]^{pp'} = \{0\}$.

Condition (G3): Suppose $\theta_i \sim_i^\Theta \theta_i'$ and $\theta_i \neq \theta_i'$. By the definition of $\sim_i^\Theta$, there exists a sequence of agent i's types $\theta_i^1, \ldots, \theta_i^K$ such that $\theta_i^1 = \theta_i$, $\theta_i^K = \theta_i'$, and $[\theta_i^k] \bowtie [\theta_i^{k+1}]$ for $k = 1, \ldots, K-1$. This implies that $[\theta_i^k]^{pp'} \bowtie [\theta_i^{k+1}]^{pp'}$ if $p \not\sim_i^X p'$ (or equivalently $p_i \neq p_i'$), which implies that θ_i and θ_i' are generalized entangled with respect to $\sim_i^X$, satisfying Condition (G3.a). Moreover, $[\theta_i]^{pp'} \cup [\theta_i']^{pp'}$ cannot be a singleton for any $p \not\sim_i^X p'$. This is because we have assumed that $v_i(\theta_i, \theta_{-i})$ is non-constant with respect to θ_{-i}, which implies that the interval $[\theta_i]$ has positive width, and consequently $[\theta_i]^{pp'}$ is not a singleton when $p_i \neq p_i'$, satisfying Condition (G3.b).

We have now established that the single-good auction environment satisfies G1D. Pick any two types θ_i, θ_i' of agent i, and any type profile $\theta_{-i} \in \Theta_{-i}$ of the other agents. First, if θ_i is higher than θ_i', then $\theta_i \succ_i^\Theta \theta_i'$, and by Proposition 3.3 we have $q(\theta_i, \theta_{-i}) \succsim_i^X q(\theta_i', \theta_{-i})$, which implies $q_i(\theta_i, \theta_{-i}) \geq q_i(\theta_i', \theta_{-i})$. Second, 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 $q_i(\theta_i, \theta_{-i}) = q_i(\theta_i', \theta_{-i})$.

A.2 Proof of Proposition 3.2

We begin with a lemma establishing that q^E is implementable if and only if it is supported by a specific transfer structure. Then we show this transfer structure exists if and only if strong single-crossing holds.

Lemma A.1. *The direct mechanism (q^E, τ) is ESP if and only if, for any agent i and any $\theta_{-i} \in \Theta_{-i}$,*

$$\tau_i(\theta_i, \theta_{-i}) = \begin{cases} \underline{p}_i(\theta_{-i}), & \text{if } q^E(\theta_i, \theta_{-i}) = 0, \\ \underline{p}_i(\theta_{-i}) + p_i(\theta_{-i}), & \text{if } q^E(\theta_i, \theta_{-i}) = 1, \end{cases}$$

for some functions $\underline{p}_i : \Theta_{-i} \to \mathbb{R}$ and $p_i : \Theta_{-i} \to \mathbb{R}$ such that

$$\max_{\tilde{\theta} \in \Theta:\ i \neq i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}) \ \leq\ p_i(\theta_{-i}) \ \leq \min_{\tilde{\theta} \in \Theta:\ i = i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}).$$

Proof. "If": Suppose transfers satisfy the stated conditions. Fix any agent i, any $\theta_i \in \Theta_i$, and any $\theta_{-i} \in \Theta_{-i}$. It suffices to verify that if agent i has true type θ_i and others report θ_{-i}, a deviation to any θ_i' that changes the allocation is unprofitable, regardless of the others' true types $\tilde{\theta}_{-i}$.

First, suppose that $q^E(\theta_i, \theta_{-i}) = 1$. In this case, the condition requiring no profitable deviation becomes

$$v_i(\theta_i, \tilde{\theta}_{-i}) - \tau_i(\theta_i, \theta_{-i}) \ \geq\ -\tau_i(\theta_i', \theta_{-i}) \quad \forall \tilde{\theta}_{-i} \in \Theta_{-i}, \tag{E1}$$

which is equivalent to

$$\min_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i, \tilde{\theta}_{-i}) \ \geq\ p_i(\theta_{-i}).$$

Since $q^E(\theta_i, \theta_{-i}) = 1$ implies that $i = i^*(\theta_i, \theta_{-i})$, it follows that

$$\min_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i, \tilde{\theta}_{-i}) \ \geq \min_{\tilde{\theta} \in \Theta:\ i = i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}) \ \geq\ p_i(\theta_{-i}),$$

where the final inequality holds by the assumption on p_i given in the lemma. Therefore, (E1) is satisfied, and the deviation is not profitable.

If instead $q^E(\theta_i, \theta_{-i}) = 0$, then an analogous argument—relying on the other inequality for p_i provided in the lemma—demonstrates that agent i cannot profit by deviating to some θ_i' such that $q^E(\theta_i', \theta_{-i}) = 1$.

"Only if": Suppose (q^E, τ) is ESP. First, observe that for any agent i, any types $\theta_i, \theta_i' \in \Theta_i$, and any profile $\theta_{-i} \in \Theta_{-i}$, if $q^E(\theta_i, \theta_{-i}) = q^E(\theta_i', \theta_{-i})$, then it must be that $\tau_i(\theta_i, \theta_{-i}) = \tau_i(\theta_i', \theta_{-i})$. Consequently, there exist functions $\underline{p}_i : \Theta_{-i} \to \mathbb{R}$ and $p_i : \Theta_{-i} \to \mathbb{R}$ such that

$$\tau_i(\theta_i, \theta_{-i}) = \begin{cases} \underline{p}_i(\theta_{-i}) & \text{if } q^E(\theta_i, \theta_{-i}) = 0, \\ \underline{p}_i(\theta_{-i}) + p_i(\theta_{-i}) & \text{if } q^E(\theta_i, \theta_{-i}) = 1. \end{cases}$$

Next, consider types $\theta_i, \theta_i' \in \Theta_i$ and a profile $\theta_{-i} \in \Theta_{-i}$ such that $q^E(\theta_i, \theta_{-i}) = 1$ and

$q^E(\theta_i', \theta_{-i}) = 0$. The ESP condition implies that for all $\tilde{\theta}_{-i} \in \Theta_{-i}$:

$$\begin{aligned} v_i(\theta_i, \tilde{\theta}_{-i}) - \tau_i(\theta_i, \theta_{-i}) &\geq -\tau_i(\theta_i', \theta_{-i}), \\ -\tau_i(\theta_i', \theta_{-i}) &\geq v_i(\theta_i', \tilde{\theta}_{-i}) - \tau_i(\theta_i, \theta_{-i}). \end{aligned}$$

Since the difference in payments is $\tau_i(\theta_i, \theta_{-i}) - \tau_i(\theta_i', \theta_{-i}) = p_i(\theta_{-i})$, these inequalities are equivalent to

$$\max_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i', \tilde{\theta}_{-i}) \leq p_i(\theta_{-i}) \leq \min_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i, \tilde{\theta}_{-i}).$$

Because this condition must hold for every choice of θ_i satisfying $q^E(\theta_i, \theta_{-i}) = 1$ (or equivalently, $i = i^*(\theta_i, \theta_{-i})$) and every θ_i' satisfying $q^E(\theta_i', \theta_{-i}) = 0$ (or equivalently, $i \neq i^*(\theta_i', \theta_{-i})$), we conclude that

$$\max_{\tilde{\theta} \in \Theta:\, i \neq i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}) \leq p_i(\theta_{-i}) \leq \min_{\tilde{\theta} \in \Theta:\, i = i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}).$$

□

Proof of Proposition 3.2. Clearly, transfer functions satisfying the conditions of Lemma A.1 exist if and only if, for any agent i and any $\theta_{-i} \in \Theta_{-i}$,

$$\max_{\tilde{\theta} \in \Theta:\, i \neq i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}) \leq \min_{\tilde{\theta} \in \Theta:\, i = i^*(\tilde{\theta}_i, \theta_{-i})} v_i(\tilde{\theta}). \tag{E2}$$

Thus, it suffices to prove that (E2) holds if and only if the environment satisfies strong single-crossing.

"If": Suppose strong single-crossing is satisfied. Fix agent i and any $\theta_{-i} \in \Theta_{-i}$. Define the set of winning types as

$$\Theta_i^*(\theta_{-i}) := \{\theta_i \in \Theta_i : i = i^*(\theta_i, \theta_{-i})\}.$$

By strong single-crossing, for any $\theta_i \in \Theta_i^*(\theta_{-i})$ and $\theta_i' \notin \Theta_i^*(\theta_{-i})$, θ_i is higher than θ_i' and the two are not entangled. Consequently,

$$\max_{\theta_i' \notin \Theta_i^*(\theta_{-i})} \max_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i', \tilde{\theta}_{-i}) \leq \min_{\theta_i \in \Theta_i^*(\theta_{-i})} \min_{\tilde{\theta}_{-i} \in \Theta_{-i}} v_i(\theta_i, \tilde{\theta}_{-i}). \tag{E3}$$

The left-hand side of (E3) is exactly the left-hand side of (E2), and the right-hand side of (E3) is exactly the right-hand side of (E2). Therefore, (E2) holds.

"Only if": Suppose (E2) holds for every agent i and every $\theta_{-i} \in \Theta_{-i}$. Then (E3) also holds. Fix $\theta_i \in \Theta_i^*(\theta_{-i})$ and $\theta_i' \notin \Theta_i^*(\theta_{-i})$. If θ_i and θ_i' were entangled, there would exist valuations for a type in $\Theta_i^*(\theta_{-i})$ that are strictly lower than valuations for a type outside $\Theta_i^*(\theta_{-i})$, contradicting (E3). Likewise, (E3) implies that it cannot be the case that θ_i' is higher than θ_i. Therefore, whenever $\theta_i \in \Theta_i^*(\theta_{-i})$, all types entangled with it, and all types higher than it, must also belong to $\Theta_i^*(\theta_{-i})$. This is exactly the strong single-crossing property.

A.3 Example 3.1

We verify that Example 3.1 satisfies G1D. Fix an agent i. For any distinct types $\theta_i, \theta_i' \in \Theta_i$, we write $\theta_i \leftrightarrow \theta_i'$ if Condition (2b) is satisfied for $\theta_i^h := \max\{\theta_i, \theta_i'\}$ and $\theta_i^l := \min\{\theta_i, \theta_i'\}$. By construction, the relation $\leftrightarrow$ is symmetric.

Define $\succsim_i^X$ on X such that $x \succsim_i^X y$ if $x \geq y$. Define $\sim_i^\Theta$ and $\succ_i^\Theta$ on Θ_i such that:

- $\theta_i \sim_i^\Theta \theta_i'$ if $\theta_i = \theta_i'$ or if there exists a finite sequence $(\theta_i^k)_{k=1}^K$ with $\theta_i^1 = \theta_i$, $\theta_i^K = \theta_i'$, such that $\theta_i^k \leftrightarrow \theta_i^{k+1}$ for every $k = 1, \ldots, K-1$;
- $\theta_i \succ_i^\Theta \theta_i'$ if $\theta_i > \theta_i'$ and $\theta_i \not\sim_i^\Theta \theta_i'$.

Clearly, $\succsim_i^X$ is complete and transitive. The argument establishing that $\succsim_i^\Theta$ is complete and transitive is analogous to that establishing the completeness and transitivity of the "weakly higher than" relation in the single-good auction model (see the proof of Proposition 3.1).

For any $x, y \in X$ such that $x > y$, we have

$$[\theta_i]^{xy} = \left[\min_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i, \theta_{-i})\, ds, \quad \max_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i, \theta_{-i})\, ds\right].$$

Suppose $\theta_i \succ_i^\Theta \theta_i'$ and $x \succ_i^X y$. Since $\theta_i \nleftrightarrow \theta_i'$, Condition (2a) implies that

$$\max_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i', \theta_{-i})\, ds \leq \min_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i, \theta_{-i})\, ds.$$

This establishes that $[\theta_i']^{xy} \to [\theta_i]^{xy}$, thereby satisfying Condition (G1). In the case where $x \sim_i^X y$ (implying $x = y$), it is immediate that $[\theta_i]^{xy} = \{0\}$, which satisfies Condition (G2).

Finally, consider $\theta_i \sim_i^\Theta \theta_i'$ with $\theta_i > \theta_i'$. If $x \not\sim_i^X y$, the set $[\theta_i]^{xy} \cup [\theta_i']^{xy}$ is not a singleton, satisfying Condition (G3.b). Furthermore, if $\theta_i \leftrightarrow \theta_i'$, Condition (2b) ensures that for any $x \succ_i^X y$,

$$\max_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i', \theta_{-i})\, ds > \min_{\theta_{-i}} \int_y^x \frac{\partial}{\partial s} u_i(s; \theta_i, \theta_{-i})\, ds.$$

This inequality rules out $[\theta_i']^{xy} \to [\theta_i]^{xy}$. Moreover, by Condition (1), for every $\hat{\theta}_{-i} \in \Theta_{-i}$ we have

$$u_i(x; \theta_i, \hat{\theta}_{-i}) - u_i(y; \theta_i, \hat{\theta}_{-i}) > u_i(x; \theta_i', \hat{\theta}_{-i}) - u_i(y; \theta_i', \hat{\theta}_{-i}),$$

so there exist $a \in [\theta_i]^{xy}$ and $b \in [\theta_i']^{xy}$ with $a > b$, which rules out $[\theta_i]^{xy} \to [\theta_i']^{xy}$. Hence $[\theta_i']^{xy} \bowtie [\theta_i]^{xy}$. Based on the construction of $\sim_i^\Theta$, this implies that if $\theta_i \sim_i^\Theta \theta_i'$, they are generalized entangled with respect to $\sim_i^X$, thus satisfying Condition (G3.a) in the definition of G1D.

A.4 Proof of Proposition 3.3

"If": Suppose q satisfies monotonicity and screenability as stipulated in the proposition. Fix any agent i and any type profile of the other agents $\hat{\theta}_{-i} \in \Theta_{-i}$. Let $\hat{X}$ denote the range of $q(\cdot, \hat{\theta}_{-i})$ as θ_i varies over Θ_i. For simplicity, assume that for any distinct $x, y \in \hat{X}$, we have either $x \succ_i^X y$ or $y \succ_i^X x$. Note that the current proof can be easily adapted

to accommodate ties (where $x \sim_i^X y$) by treating such alternatives as a single equivalent alternative.

Index the alternatives in $\hat{X}$ as $x^1, \ldots, x^l$ such that $x^l \succ_i^X x^{l-1} \succ_i^X \cdots \succ_i^X x^1$. For each $k = 1, \ldots, l$, let $\Theta_i^k := \{\theta_i \in \Theta_i : q(\theta_i, \hat{\theta}_{-i}) = x^k\}$. Furthermore, for $k = 2, \ldots, l$, define

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

By monotonicity and screenability, if $\theta_i \in \Theta_i^k$ and $\theta_i' \in \Theta_i^{k'}$ with $k > k'$ (or equivalently, $q(\theta_i, \hat{\theta}_{-i}) \succ_i^X q(\theta_i', \hat{\theta}_{-i})$), then $\theta_i \succ_i^\Theta \theta_i'$. This implies that $[\theta_i']^{x^k x^{k-1}} \to [\theta_i]^{x^k x^{k-1}}$. Thus, for any $k = 2, \ldots, l$, 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 $\hat{\theta}_{-i}$, construct the transfers for agent i as follows:

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

We now show that if the other agents report $\hat{\theta}_{-i}$, truthful reporting is optimal for agent i regardless of the actual types of the other agents. Suppose agent i has type θ_i. Thus, $\theta_i \in \Theta_i^k$ for some $k \in \{1, \ldots, l\}$. Let the true types of the other agents be θ_{-i}.

Observe that agent i has no incentive to report any $\theta_i' \in \Theta_i^k$, as this changes neither the chosen alternative nor the transfer payment. If she deviates to a report $\theta_i' \in \Theta_i^{k'}$ 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). Similarly, deviating to any $\theta_i' \in \Theta_i^{k'}$ where $k' < k$ is unprofitable by inequality (I2). Since $\hat{\theta}_{-i}$ was arbitrarily chosen, it follows that q is implementable.

"Only if": Suppose q is implementable with supporting transfer function $(\tau_1, \ldots, \tau_n)$. Fix any agent i and any $\hat{\theta}_{-i} \in \Theta_{-i}$. Pick a pair of agent i's types θ_i and θ_i'. Denote $x := q(\theta_i, \hat{\theta}_{-i})$ and $y := q(\theta_i', \hat{\theta}_{-i})$. Incentive compatibility for agent i of type θ_i requires

$$u_i(x; \theta_i, \theta_{-i}) - \tau_i(\theta_i, \hat{\theta}_{-i}) \geq u_i(y; \theta_i, \theta_{-i}) - \tau_i(\theta_i', \hat{\theta}_{-i}) \quad \forall \theta_{-i} \in \Theta_{-i},$$

which implies

$$\min_{\theta_{-i}} u_i^{xy}(\theta_i, \theta_{-i}) \geq \tau_i(\theta_i, \hat{\theta}_{-i}) - \tau_i(\theta_i', \hat{\theta}_{-i}).$$

Similarly, incentive compatibility for agent i of type θ_i' implies

$$\tau_i(\theta_i, \hat{\theta}_{-i}) - \tau_i(\theta_i', \hat{\theta}_{-i}) \geq \max_{\theta_{-i}} u_i^{xy}(\theta_i', \theta_{-i}).$$

Combining these inequalities yields

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

To show monotonicity, suppose $\theta_i \succ_i^{\Theta} \theta_i'$, and we wish to show that $x \succsim_i^X y$. Observe that $y \not\succ_i^X x$. If it were the case that $y \succ_i^X x$, we would have $[\theta_i']^{yx} \to [\theta_i]^{yx}$. This implies

$$\max_{\theta_{-i}} u_i^{yx}(\theta_i, \theta_{-i}) > \min_{\theta_{-i}} u_i^{yx}(\theta_i', \theta_{-i}),$$

which contradicts Inequality (I3). Thus, we must 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 (G3.a) implies that there exists a sequence of types $\theta_i^1, \ldots, \theta_i^K$ such that $\theta_i^1 = \theta_i$, $\theta_i^K = \theta_i'$, and $[\theta_i^k]^{xy} \bowtie [\theta_i^{k+1}]^{xy}$ for $k = 1, \ldots, K-1$. Note that for each $k = 1, \ldots, K-1$, we have $\theta_i^k \sim_i^{\Theta} \theta_i^{k+1}$. Otherwise, by Condition (G1), we would have either $[\theta_i^k]^{xy} \to [\theta_i^{k+1}]^{xy}$ or $[\theta_i^{k+1}]^{xy} \to [\theta_i^k]^{xy}$ (since $x \not\sim_i^X y$).

Fix any $k \in \{1, \ldots, K-1\}$. Denote $w := q(\theta_i^k, \hat{\theta}_{-i})$ and $z := q(\theta_i^{k+1}, \hat{\theta}_{-i})$. We want to show that $w \sim_i^X z$. Suppose $w \not\sim_i^X z$; then by Condition (G3.a), $[\theta_i^{k+1}]^{wz} \bowtie [\theta_i^k]^{wz}$. Now consider the incentive compatibility constraints:

- For agent i of type θ_i^k not to deviate to reporting type θ_i^{k+1}:

$$\min_{\theta_{-i}} u_i^{wz}(\theta_i^k, \theta_{-i}) \geq \tau_i(\theta_i^k, \hat{\theta}_{-i}) - \tau_i(\theta_i^{k+1}, \hat{\theta}_{-i}).$$

- For agent i of type θ_i^{k+1} not to deviate to reporting type θ_i^k:

$$\tau_i(\theta_i^k, \hat{\theta}_{-i}) - \tau_i(\theta_i^{k+1}, \hat{\theta}_{-i}) \geq \max_{\theta_{-i}} u_i^{wz}(\theta_i^{k+1}, \theta_{-i}).$$

Combining these, we derive:

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

Recall that there are three possibilities (see page 14) for $[\theta_i^{k+1}]^{wz} \bowtie [\theta_i^k]^{wz}$. Inequality (I4) leads to the third possibility: the two intervals coincide as degenerate intervals, which violates Condition (G3.b). Thus, we must have $w \sim_i^X z$. Since k was arbitrarily

chosen from $\{1, \ldots, K-1\}$, we conclude that

$$x = q(\theta_i^1, \hat{\theta}_{-i}) \sim_i^X q(\theta_i^2, \hat{\theta}_{-i}) \sim_i^X \cdots \sim_i^X q(\theta_i^K, \hat{\theta}_{-i}) = y,$$

which contradicts the assumption that $x \not\sim_i^X y$. It follows that $x \sim_i^X y$.

A.5 Proof of Propositions 3.4

Given any agent i, any finite sequence of her types $\boldsymbol{\theta_i} = (\theta_i^1, \ldots, \theta_i^K)$, and any $\theta'_{-i} \in \Theta_{-i}$, define

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

"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 for agent i that start with $\tilde{\theta}_i$ and terminate with θ_i.

For any $\theta_i \in \Theta_i$ and any $\theta'_{-i} \in \Theta_{-i}$, define

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

Observe that GCM implies $V(\tilde{\theta}_i, \theta'_{-i}) \leq 0$. Moreover, since $G(\boldsymbol{\theta_i}, \theta'_{-i}) = 0$ for the trivial sequence $\boldsymbol{\theta_i} = (\tilde{\theta}_i, \tilde{\theta}_i)$, we have $V(\tilde{\theta}_i, \theta'_{-i}) = 0$ for any $\theta'_{-i} \in \Theta_{-i}$.

Notice that

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

Since $V(\tilde{\theta}_i, \theta'_{-i}) = 0$, it follows that

$$V(\theta_i, \theta'_{-i}) \leq -\max_{\theta_{-i}} \Big\{ u_i\Big(q(\tilde{\theta}_i, \theta'_{-i}); \theta_i, \theta_{-i}\Big) - u_i\Big(q(\theta_i, \theta'_{-i}); \theta_i, \theta_{-i}\Big) \Big\}.$$

Thus, $V(\cdot, \theta'_{-i})$ is bounded from above and well-defined.

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

and $\theta_{-i}, \theta'_{-i} \in \Theta_{-i}$:

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

It follows that q is implementable.

"Only if": Suppose q is implementable with supporting transfer function $(\tau_1, \ldots, \tau_n)$. Fix any agent i and any sequence $(\theta_i^1, \ldots, \theta_i^K)$ of types for agent i where $\theta_i^K = \theta_i^1$. Incentive compatibility implies that for any $k = 1, \ldots, K-1$ and any $\theta'_{-i} \in \Theta_{-i}$:

$$\forall \theta_{-i} \in \Theta_{-i}: \quad u_i\Big(q(\theta_i^{k+1}, \theta'_{-i}); \theta_i^k, \theta_{-i}\Big) - \tau_i(\theta_i^{k+1}, \theta'_{-i}) \leq u_i\Big(q(\theta_i^k, \theta'_{-i}); \theta_i^k, \theta_{-i}\Big) - \tau_i(\theta_i^k, \theta'_{-i}),$$

which in turn implies that

$$\max_{\theta_{-i}} \Big\{u_i\Big(q(\theta_i^{k+1}, \theta'_{-i}); \theta_i^k, \theta_{-i}\Big) - u_i\Big(q(\theta_i^k, \theta'_{-i}); \theta_i^k, \theta_{-i}\Big)\Big\} \leq \tau_i(\theta_i^{k+1}, \theta'_{-i}) - \tau_i(\theta_i^k, \theta'_{-i}). \tag{I_k}$$

Summing the inequalities $I_1, \ldots, I_{K-1}$, the terms on the right-hand side form a telescoping sum that vanishes (since $\theta_i^K = \theta_i^1$), leaving:

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

It follows that q satisfies GCM.

A.6 Proof of Proposition 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 $G(\boldsymbol{\theta_i}, \theta'_{-i})$ from the proof of Proposition 3.4.

Suppose q satisfies GCM. Fix any agent i and any $\theta'_{-i} \in \Theta_{-i}$. Let μ be the type-alternative matching for agent i induced by q given θ'_{-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.[17] That is, 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) $\theta_i^{k+1} = \pi(\theta_i^k)$ for

[17] See, for example, Theorem 9.8 of Fraleigh (2002), p. 89.

every $k = 1, ..., K-2$. 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 $W(p) := G(\boldsymbol{\theta_i}, \theta'_{-i})$.

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

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

Thus, μ is reshuffling-efficient.

Now suppose that for any agent i and any $\theta'_{-i} \in \Theta_{-i}$ from the other agents, the type-alternative matching μ for agent i induced by q given θ'_{-i} is reshuffling-efficient. However, suppose GCM is not satisfied. Then, for some θ'_{-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 $G(\boldsymbol{\theta_i}, \theta'_{-i}) > 0$. Let $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}} = (\hat{\theta}_i^1, ..., \hat{\theta}_i^K)$ be the *shortest* such sequence given θ'_{-i}.

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

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

It is straightforward to verify that π is a permutation, as each type appearing in the sequence $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$ 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' be the reshuffling obtained from π, i.e., $\mu'(\theta_i) = \mu(\pi(\theta_i))$. Thus we have

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

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 $\sum_{k=a}^{b-1} \max_{\theta_{-i}} \Big\{ u_i\Big(q(\hat{\theta}_i^{k+1}, \theta'_{-i}); \hat{\theta}_i^k, \theta_{-i}\Big) - u_i\Big(q(\hat{\theta}_i^k, \theta'_{-i}); \hat{\theta}_i^k, \theta_{-i}\Big) \Big\}$ 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

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

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 $G(\cdot, \theta'_{-i})$ is positive, because $\tilde{\boldsymbol{\theta}}_{\boldsymbol{i}}$ is by construction strictly shorter than $\hat{\boldsymbol{\theta}}_{\boldsymbol{i}}$. This concludes the proof.

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.

Chen, Yan and Onur Kesten (2017), "Chinese college admissions and school choice reforms: A theoretical analysis." *Journal of Political Economy*, 125, 99–139.

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 (2025), "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.

Vohra, Rakesh V (2011), *Mechanism design: a linear programming approach*, volume 47. Cambridge University Press.

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