Knowledge Transfer and Strategic Similarity*

Joseph Feffer Filip Tokarski
Stanford GSB Stanford GSB

Abstract

This paper studies when strategic understanding acquired in one mechanism can be transferred to another. We introduce a framework in which agents' knowledge is represented as a set of payoff comparisons they can make, and use it to formalize what it means to understand that a strategy profile is an equilibrium. We first apply this framework to mechanisms that are *strategically equivalent*—that is, share the same game form up to relabeling of actions—and show that agents' understanding of equilibrium transfers across such mechanisms once the relevant action correspondences are explained to them. We then define *strategic similarity*, a weaker notion that allows not only actions but also types to be remapped, and show that understanding of equilibrium transfers across strategically similar mechanisms once agents recognize how actions and types correspond. Applications include single-item auctions, scoring auctions, and nonlinear pricing with capacity constraints.

1 Introduction

Strategizing in new economic settings can be difficult even for sophisticated agents. However, when a new setting is structurally related to one that agents already understand, explaining this relationship can let agents reuse their existing strategic reasoning in it. Consider, for instance, the case where two mechanisms are *strategically equivalent*, in the sense that they share the same game form and differ only in how choices are presented. A standard example is the sealed-bid first-price auction and the Dutch auction, in which a price clock descends until some bidder stops it and wins the object at the current price (Vickrey, 1961; Milgrom and Weber, 1982). Although the two formats frame choices differently, each effectively asks a bidder to choose the price at which she is willing to buy: in the first-price auction she names this price directly as her bid, while in the Dutch auction she waits for the clock to reach it and then stops it. A bidder who knows how to play in one and recognizes how the two representations relate can apply that knowledge to the other.

*We thank Mohammad Akbarpour, Roberto Corrao, Laura Doval, Piotr Dworczak, Paul Milgrom, Ilya Segal, Andrzej Skrzypacz, Takuo Sugaya and Eric Tang for helpful comments.

The transfer of strategic reasoning, however, need not be limited to cases in which two mechanisms are just alternative representations of the same game. Some genuinely distinct mechanisms still require agents to reason through analogous tradeoffs and apply the same strategic principles. Consider, for example, an agent who understands how to play in a first-price auction with a specific reserve price. To bid optimally, she must be able to weigh the probability of winning against her conditional expected surplus, translate her beliefs about others' values into beliefs about the bids she will face, and understand how the reserve affects participation. Now, suppose that same agent were asked to play in a first-price auction with a different reserve. If she recognized that the two mechanisms share much of their strategic structure, she could adapt her understanding of the original auction to the new one. In this paper, we develop a notion of *strategic similarity* that generalizes this idea.

Understanding what mechanisms are strategically similar is practically important in settings where the same agent interacts with many versions of the same underlying problem. An example of such a setting comes from the Chilean public procurement system for pharmaceuticals studied by Allende et al. (2024). There, hospitals purchase supplies by holding procurement auctions on a public platform: each auction specifies a scoring rule that aggregates price and other product attributes into a single score, which is then used to determine the winner. Crucially, hospitals can express their preferences over price, delivery terms, and other product characteristics by customizing the weights the scoring rule puts on them. This also means, however, that manufacturers and wholesalers may bid in hundreds of auctions with potentially different strategic properties. In order to mitigate this problem, the platform fixes a baseline format and restricts hospitals to fine-tuning its parameters. However, parameter changes that adjust a mechanism's rules in a simple way may nevertheless alter its strategic properties. To minimize the need for bidders to repeatedly re-learn how to play, it is therefore important to understand which templates and parameter choices preserve the underlying strategic logic, and which ones generate genuinely different strategic problems.

This paper makes two main contributions. First, it develops a theory of reasoning transfer across strategic settings. To formalize this phenomenon, we model an agent's understanding of mechanisms through a set of payoff comparisons she can make, as well as her beliefs about which comparisons others are aware of. We use this framework to define when a strategy profile is commonly known to be an equilibrium. We then model agents' ability to transfer knowledge across mechanisms by specifying payoff equivalences from which they can derive new comparisons. This lets us characterize when understanding similarities in mechanisms' rules allows agents to carry over their understanding of their equilibria. Proposition 1 formalizes this idea in the context of strategic equivalence.

Our second contribution is the notion of *strategic similarity*, which generalizes the idea of strategic equivalence. Intuitively, two mechanisms are strategically similar if, after appropriately relabeling actions and remapping types, the payoffs from any action profile in one mechanism correspond with the payoffs from the analogous profile in the other. Allowing types to be remapped alongside actions provides a sense in which two mechanisms may share the same strategic structure even when they do not share the same game form. Proposition 2 then shows that, once the relevant correspondences between types and actions are explained, agents can transfer their understanding of equilibria across strategically similar mechanisms.

We apply our framework to a range of economic settings. In the context of single-item auctions, we show that kth-price auctions with reserves form a maximal class of strategically similar mechanisms within a natural set of auctions. This aligns with common practice in auction houses and ad exchanges, which typically commit to a fixed format and use the reserve price as their main tuning parameter. We further show that a single strategically similar class can implement the revenue-optimal auction for every symmetric regular prior, but not for every symmetric prior. The next application studies scoring auctions, commonly used in procurement, where a buyer awards a contract based on a score that aggregates price and quality. We show that varying the weight on price preserves strategic similarity under a linear scoring rule, but not under a rule in which the price score depends on a bidder's price relative to the lowest competing price. Finally, we study non-linear pricing of services produced using scarce capacity. We show that charging for the underlying input can preserve strategic similarity across changes in productivity in cases where charging for realized output does not.

The rest of the paper is organized as follows. Section 2 discusses related literature. Section 3 presents the model and introduces the concepts of payoff comparisons, knowledge, and knowledge transfer. Section 4 studies strategic equivalence and shows how explaining action correspondences lets agents transfer knowledge of equilibria across equivalent mechanisms. Section 5 defines strategic similarity, establishes a knowledge transfer result, and applies the concept to single-item auctions, scoring auctions, and input- versus output-based pricing.

2 Related literature

Our paper builds on experimental work on reasoning transfer in strategic environments. Rick and Weber (2010) show that when subjects learn a general solution method such as iterated dominance, they are often able to apply it in new but structurally related environments. Similarly, Cooper and Kagel (2008) find that experience in one signaling game improves equilibrium-consistent play in related games. In auction settings, Harstad (2000) documents partial transfer of experience across related auction formats. Breitmoser and Schweighofer-Kodritsch (2022) find that performance in a sealed-bid second-price auction improves substantially when subjects are shown a clock implementation that yields the same outcome. According to the authors, subjects find the dynamic setting easier to navigate, and showing them the clock implementation makes the analogy between the sealed-bid auction and the dynamic format more salient. This, in turn, helps participants transfer their strategic understanding to the sealed-bid format.

The effort to model agents' reasoning about mechanisms also connects to the epistemic game theory literature, which formalizes agents' knowledge and uses it to justify solution concepts. The rationalizability literature characterizes strategies consistent with common knowledge of rationality (Bernheim, 1984; Pearce, 1984), while Aumann and Brandenburger (1995) identify epistemic conditions under which conjectures constitute a Nash equilibrium. We use similar tools to ask a distinct question: we assume agents can identify an equilibrium in one mechanism and ask under what conditions this knowledge lets them identify a corresponding equilibrium in another. Modeling agents' limited understanding of payoff comparisons also relates to papers studying imperfect knowledge of the game structure: Feinberg (2021) models unawareness in games through a collection of games representing players' perceptions of the strategic situation. Copic and Galeotti (2006) study normal-form games in which parts of the game are not common knowledge and define an awareness equilibrium.

Our notion of strategic similarity also connects to the study of game isomorphisms and payoff transformations. McKinsey (2016) asks when games share the same strategic structure up to relabeling, and Elmes and Reny (1994) extend related ideas to extensive-form games. Moulin and Vial (1978) define strategic equivalence by preservation of preference rankings over mixed strategies, while Morris and Ui (2005) and Tewolde and Conitzer (2021) study payoff transformations that preserve best responses or Nash equilibria. Our notion extends these ideas to incomplete-information settings, which lets us apply it beyond matrix games to mechanism design environments.

Another related strand of work models reasoning by analogy in decision theory and games. Gilboa and Schmeidler (1995) propose a model of choice under uncertainty in which agents look to similar past cases when choosing among actions. Samuelson (2001) similarly treats analogies as a way of economizing on scarce reasoning resources: agents maintain a costly stock of models and apply the one that appears most suitable to the strategic interaction at hand. Jehiel (2005) develops the concept of analogy-based expectation equilibrium, in which agents bundle contingencies into analogy classes when forming expectations and best respond to those simplified representations. While these papers treat analogical reasoning as a form of bounded rationality, we instead study agents who respond optimally but can more easily identify the right strategy in mechanisms that are strategically close to ones they already understand.

Lastly, this paper connects to a literature studying mechanisms that are simple for participants to play. The practical importance of simplicity is emphasized by Li (2024), who highlights that desirable theoretical properties of mechanisms can break down when agents fail to recognize incentives and respond to them correctly. Various approaches have been proposed to conceptualize such cognitive limitations. Li (2017) ask when equilibrium play is transparent to agents with limited contingent reasoning. Pycia and Troyan (2023) develop a theory of simplicity in extensive-form environments with limited planning horizons. Börgers and Li (2019) focus on belief sophistication, defining strategically simple mechanisms as those in which an agent can identify an optimal action using only first-order beliefs about others' preferences. While this literature studies how difficult it is to play a particular mechanism, we study how difficult it is to play within a *class* of mechanisms, given that the agent already understands one of them. Our perspective is therefore complementary: in many applications, mechanisms should not only be simple in isolation, but their strategic logic should also be stable across related instances, so that agents do not need to re-learn how to play each time the environment changes. In this sense, our perspective is aligned with Brooks and Du (2025), who argue that practically useful mechanisms should be portable across a range of settings.

3 Model

3.1 Environments, mechanisms, and equilibria

We focus on the similarity of mechanisms designed for the same kind of economic problem. To make this precise, we formally define an *environment* as a tuple $(\mathcal{I}, \mathcal{Y}, \mathcal{T}, \mathcal{U})$, where $\mathcal{I}$ is a finite nonempty set of agents, $\mathcal{Y}$ is an outcome space, $\mathcal{T} = \prod_{i\in\mathcal{I}} \mathcal{T}_i$ is the type space, and $\mathcal{U} = (u_i)_{i\in\mathcal{I}}$ is the profile of Bernoulli utility functions giving the payoff each type of every agent gets from every outcome: $u_i : \mathcal{T}_i \times \mathcal{Y} \to \mathbb{R}$.[1] Agents maximize expected utility.

In each such environment, many mechanisms may be used to select an outcome based on agents' actions. We define a *mechanism* for an environment $(\mathcal{I}, \mathcal{Y}, \mathcal{T}, \mathcal{U})$ as a pair $X = (\mathcal{A}, \Phi)$, where $\mathcal{A} = \prod_{i\in\mathcal{I}} \mathcal{A}_i$ specifies agents' action spaces and $\Phi : \mathcal{A} \to \Delta(\mathcal{Y})$ is the outcome rule, mapping each action profile to a lottery over outcomes. When $\Phi(a)$ is degenerate at some outcome $y \in \mathcal{Y}$, we abuse notation and write $\Phi(a) = y$. For $i \in \mathcal{I}$, we also write

$$\mathcal{A}_{-i} := \prod_{j\in\mathcal{I}\setminus\{i\}} \mathcal{A}_j, \qquad \mathcal{T}_{-i} := \prod_{j\in\mathcal{I}\setminus\{i\}} \mathcal{T}_j.$$

We illustrate the distinction between environments and mechanisms with an example.

Example 1. *Consider an environment in which a single good can be allocated to one of two agents, or remain unassigned. Each agent i has value $t_i \in \mathbb{R}$ for the good. This environment can be written as*

$$\left(\mathcal{I}, \mathcal{Y}, \mathbb{R}^{\mathcal{I}}, (u_i)_{i\in\mathcal{I}}\right), \qquad \text{where} \qquad \mathcal{I} = \{1, 2\}, \qquad \mathcal{Y} := \mathcal{I} \cup \{\varnothing\}, \qquad u_i(t_i, y) = t_i \mathbb{1}\{y = i\}.$$

We now present two mechanisms for this environment. In the first mechanism, each agent reports whether she wants the good. If exactly one agent says yes, she gets it; otherwise the good is discarded:

$$X_1 = \left(\mathcal{A}_1^1 \times \mathcal{A}_2^1, \Phi^1\right), \qquad \text{where} \qquad \mathcal{A}_i^1 = \{yes, no\}, \quad \Phi^1(a_1, a_2) = \begin{cases} 1 & \text{if } (a_1, a_2) = (yes, no), \\ 2 & \text{if } (a_1, a_2) = (no, yes), \\ \varnothing & \text{otherwise.} \end{cases}$$

In the second mechanism, agents have no choice and the good is assigned uniformly at random:

$$X_2 = \left(\mathcal{A}_1^2 \times \mathcal{A}_2^2, \Phi^2\right), \qquad \text{where} \qquad \mathcal{A}_i^2 = \{*\}, \qquad \Phi^2(*, *) = \tfrac{1}{2}\delta_1 + \tfrac{1}{2}\delta_2.$$

[1] Note that types $t_i \in \mathcal{T}_i$ are *payoff types*: they specify preferences over outcomes, rather than information types or hierarchies of beliefs in the sense of Harsanyi (1967) and Mertens and Zamir (1985). Thus, the assumption that an agent's utility depends only on her own type is substantive.

We now define the *Bayes–Nash equilibrium* of a mechanism. For any mechanism $X = (\mathcal{A}, \Phi)$, let

$$U_j^X[t_j, A] := \mathbb{E}_{m \sim A}\left[\int_{\mathcal{Y}} u_j(t_j, y)\, d\Phi(m)(y)\right]$$

denote the expected utility agent j of type t_j gets when the mixed action profile $A \in \Delta(\mathcal{A})$ is played. For a pure action profile $a \in \mathcal{A}$, we abuse notation and write

$$U_j^X[t_j, a] := U_j^X[t_j, \delta_a].$$

Now, let a *strategy profile* $\sigma = (\sigma_i)_{i \in \mathcal{I}}$ be a collection of maps $\sigma_i : \mathcal{T}_i \to \Delta(\mathcal{A}_i)$ specifying a lottery over available actions for each type of agent i. Given a strategy profile σ and a prior $F \in \Delta(\mathcal{T})$ over all agents' types, let $A_{-i}(\sigma, F \mid t_i) \in \Delta(\mathcal{A}_{-i})$ denote the induced distribution over other agents' actions, given σ and F, conditional on agent i's type being t_i.

Definition 1. *A strategy profile σ is a **Bayes–Nash equilibrium** of mechanism $X = (\mathcal{A}, \Phi)$ under prior F if for all $i \in \mathcal{I}$ and $t_i \in \mathcal{T}_i$,*

$$\sigma_i(t_i) \in \arg\max_{A_i \in \Delta(\mathcal{A}_i)} U_i^X\big[t_i,\ A_i \otimes A_{-i}(\sigma, F \mid t_i)\big].$$

Here $\otimes$ forms the product distribution: $A_j \otimes A_{-j} \in \Delta(\mathcal{A})$ is the mixed action profile induced by independent randomization according to $A_j \in \Delta(\mathcal{A}_j)$ and $A_{-j} \in \Delta(\mathcal{A}_{-j})$.

3.2 Knowledge of payoff comparisons

We now fix an environment $(\mathcal{I}, \mathcal{Y}, \mathcal{T}, \mathcal{U})$ and consider agents' understanding of the payoff structures of mechanisms designed for it. We express this understanding using *payoff comparisons*, which are statements of the form

$$U_j^X[t_j, A] \geq U_j^X[t_j, A'],$$

where $X = (\mathcal{A}, \Phi)$ is a mechanism, $j \in \mathcal{I}$ is an agent, $t_j \in \mathcal{T}_j$ is a type, and $A, A' \in \Delta(\mathcal{A})$ are mixed action profiles. Intuitively, this payoff comparison says that agent j of type t_j gets a higher payoff when action profile A is played than she gets when A' is played. We use $\mathfrak{R}$ to denote the set of all possible payoff comparisons.

Remark 1. *When agents randomize their actions independently, A_{-j} can be interpreted as agent j's belief about how others play, while A_j is the mixed action she is evaluating. Under this interpretation, a payoff comparison says that agent j of type t_j weakly prefers mixed action A_j when her belief about others' actions is A_{-j} to mixed action A'_j when her belief about others' actions is A'_{-j}.*

Remark 2. *We emphasize that payoff comparisons capture an agent's knowledge of how different action profiles rank within a given mechanism, not her knowledge of how she values outcomes directly. In* *particular, consider two action profiles in mechanism X and two action profiles in mechanism X′ such that the corresponding pairs induce the same outcomes. Our framework allows the agent to know which of the two profiles is better in X while being unaware of which is better in X′. This reflects the fact that recognizing the consequences of actions requires understanding how the mechanism maps actions to outcomes.*

We capture which payoff comparisons each agent is aware of through an *awareness profile* $\omega \in \Omega := \prod_{i \in \mathcal{I}} 2^{\mathfrak{R}}$:

$$\omega = (\mathcal{R}_i^{\omega})_{i \in \mathcal{I}}, \qquad \mathcal{R}_i^{\omega} \subseteq \mathfrak{R}.$$

Here $\mathcal{R}_i^{\omega}$ is the set of payoff comparisons agent i is aware of under ω. We also model an agent's knowledge of which comparisons others can make; it is captured by a nonempty *knowledge set*

$$\mathcal{K}_i \subseteq \Omega,$$

consisting of the awareness profiles that agent i considers possible. Following the standard possibility-set formulation of knowledge, we say that agent i *knows* an event $E \subseteq \Omega$ if $\mathcal{K}_i \subseteq E$ (Aumann, 1976). For illustration, consider the following example.

Example 2. *Consider the environment from Example 1. Suppose agent $k \in \mathcal{I}$ knows that agent 1 is aware that, in mechanism X_1, agent 1 is always better off saying "yes" whenever her type is positive. In the notation above, this means*

$$\mathcal{K}_k \subseteq \bigcap_{t_1 \geq 0,\ a_2 \in \mathcal{A}_2^1} \left\{ \omega \in \Omega : \left\{ U_1^{X_1}[t_1, (yes, a_2)] \geq U_1^{X_1}[t_1, (no, a_2)] \right\} \in \mathcal{R}_1^{\omega} \right\}.$$

That is, agent k only entertains awareness profiles in which agent 1 is aware of each of these comparisons.

We say that an event $E \subseteq \Omega$ is *common knowledge* if $\mathcal{K}_i \subseteq E$ for every $i \in \mathcal{I}$. We also say that a comparison $r \in \mathfrak{R}$ is common knowledge if, for every $i \in \mathcal{I}$,

$$\mathcal{K}_i \subseteq \bigcap_{j \in \mathcal{I}} \{\omega \in \Omega : r \in \mathcal{R}_j^{\omega}\}.$$

Remark 3 (Relation to Aumann (1976))**.** *Our knowledge sets $(\mathcal{K}_i)_{i \in \mathcal{I}}$ play the same structural role as information partitions in the standard model of Aumann (1976): they are primitives of the epistemic structure, not objects about which agents are themselves uncertain. The difference is that an Aumann partition specifies, for every possible state ω, the set of states agent i considers possible when ω is the true state. Here we only keep the possibility set that is relevant for reasoning about the actual awareness profile. Thus, each agent is assigned a single set $\mathcal{K}_i \subseteq \Omega$, rather than a state-dependent information cell $\Pi_i(\omega)$. Also, since knowledge sets are primitives of the epistemic structure, the condition that $\mathcal{K}_i \subseteq E$ for every i already captures the relevant higher-order knowledge. If $\mathcal{K}_i \subseteq E$ for every i, then every agent knows E. Since the knowledge sets themselves are fixed features of the epistemic structure, agents also know that others know E, know that others know that others know E, and so on.*

Using this framework, we can define what it means for it to be common knowledge that a strategy profile σ is a Bayes–Nash equilibrium.

Definition 2. *Fix a mechanism $X = (\mathcal{A}, \Phi)$, a prior $F \in \Delta(\mathcal{T})$, and an equilibrium σ of X under F. We say that it is **common knowledge that** σ **is an equilibrium of** X **at** F if for each agent $j \in \mathcal{I}$, type $t_j \in \mathcal{T}_j$, and mixed deviation $A_j \in \Delta(\mathcal{A}_j)$, the comparisons*

$$U_j^X\left[t_j, \sigma_j(t_j) \otimes A_{-j}(\sigma, F \mid t_j)\right] \geq U_j^X\left[t_j, A_j \otimes A_{-j}(\sigma, F \mid t_j)\right],$$

are common knowledge. Here $A_{-j}(\sigma, F \mid t_j)$ is the distribution over opponents' actions induced by σ_{-j} and the conditional distribution of t_{-j} given t_j.

Intuitively, it is common knowledge that σ is an equilibrium if all agents know that, under this strategy profile, no agent can gain by deviating unilaterally.

3.3 Knowledge transfer

We now formalize how agents who understand certain correspondences between mechanisms can transfer payoff comparisons between them. To model this form of reasoning, we first define *payoff equivalences*, which are statements of the form

$$(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t_j'),$$

where $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ are mechanisms, $j \in \mathcal{I}$, $t_j, t_j' \in \mathcal{T}_j$, $A \in \Delta(\mathcal{A})$, and $A' \in \Delta(\mathcal{A}')$. We use $\mathfrak{E}$ to denote the set of all payoff equivalences. Intuitively, this statement says that, for agent j, the payoff from A under type t_j in mechanism X corresponds to the payoff from A' under type t_j' in mechanism X'. The exact meaning of a payoff equivalence comes from how agents can use it to extend the comparisons they already know. To capture this, we define what it means for a comparison set $\mathcal{R}$ to be extended through a given equivalence set:

Definition 3. *A set of comparisons $\mathcal{R} \subset \mathfrak{R}$ is **extended through an equivalence set** $\mathcal{E} \subseteq \mathfrak{E}$ if the pair $(\mathcal{R}, \hat{\mathcal{E}})$ satisfies the following properties for some $\hat{\mathcal{E}} \supseteq \mathcal{E}$:*

1. ***Transfer.*** *Suppose $\left\{(X, A^1, t_j) \underset{j}{\Leftrightarrow} (X', A'^1, t_j'),\ (X, A^2, t_j) \underset{j}{\Leftrightarrow} (X', A'^2, t_j')\right\} \subseteq \hat{\mathcal{E}}$. Then:*

$$\left\{U_j^X[t_j, A^1] \geq U_j^X[t_j, A^2]\right\} \in \mathcal{R} \quad \Longrightarrow \quad \left\{U_j^{X'}[t_j', A'^1] \geq U_j^{X'}[t_j', A'^2]\right\} \in \mathcal{R}.$$

2. ***Symmetry.*** *$\left\{(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t_j')\right\} \in \hat{\mathcal{E}}$ if and only if $\left\{(X', A', t_j') \underset{j}{\Leftrightarrow} (X, A, t_j)\right\} \in \hat{\mathcal{E}}$.*

3. ***Mixtures.*** *Let (S,μ) be a probability space and let $A^s \in \Delta(\mathcal{A})$, $A'^s \in \Delta(\mathcal{A}')$ be measurable in s. If $\{(X, A^s, t_j) \underset{j}{\Leftrightarrow} (X', A'^s, t'_j)\} \in \hat{\mathcal{E}}$ for every $s \in S$, then*

$$\left\{ \left(X,\, \textstyle\int_S A^s \, d\mu,\, t_j\right) \underset{j}{\Leftrightarrow} \left(X',\, \textstyle\int_S A'^s \, d\mu,\, t'_j\right) \right\} \in \hat{\mathcal{E}}.$$

Intuitively, a set of comparisons $\mathcal{R}$ is extended through $\mathcal{E}$ if it contains all the inferences the agent could make using those equivalences. More specifically, an agent can reason with $\mathcal{E}$ in two ways. First, she can use these equivalences to transfer knowledge between mechanisms. This is captured by the transfer property: once an agent knows two payoff situations are equivalent, she can carry payoff comparisons from one to the other. Second, she can extend the set of equivalences themselves by reversing them (the symmetry property) and by combining them across mixtures over action profiles (the mixtures property). The latter is important because it ensures that if equivalence holds pointwise across action profiles, then it also holds after randomizing over them; this lets agents transfer knowledge about mixed actions.

Importantly, different collections of mechanisms will admit different sets of equivalences that agents can safely use, in the sense that these equivalences will never generate false comparisons from true ones. We capture this with the following definition:

Definition 4. *We say that a set of equivalences $\mathcal{E} \subseteq \mathfrak{E}$ is* ***valid*** *if for every $\mathcal{R} \subseteq \mathfrak{R}$ containing only true comparisons, the smallest $\hat{\mathcal{R}} \supseteq \mathcal{R}$ that is extended through $\mathcal{E}$ also contains only true comparisons.*[2]

As it turns out, valid equivalence sets admit a clean characterization.

Theorem 1. *An equivalence set $\mathcal{E}$ is valid if and only if, for every agent j, every two mechanisms $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$, and every two types $t_j, t'_j \in \mathcal{T}_j$, there exists a positive affine function*

$$\ell^j_{(X,t_j)\to(X',t'_j)} : \mathbb{R} \to \mathbb{R} \quad \text{such that} \quad \ell^j_{(X',t'_j)\to(X,t_j)} = \left(\ell^j_{(X,t_j)\to(X',t'_j)}\right)^{-1}, \tag{1}$$

and such that for all $A \in \Delta(\mathcal{A})$ and $A' \in \Delta(\mathcal{A}')$,

$$\left\{(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t'_j)\right\} \in \mathcal{E} \quad \Longrightarrow \quad U_j^{X'}[t'_j, A'] = \ell^j_{(X,t_j)\to(X',t'_j)}\left(U_j^X[t_j, A]\right).$$

The affine structure arises from requiring equivalences to remain valid when combined through mixtures. Since expected payoffs under mixtures are convex combinations of the payoffs being mixed, the transformation relating equivalent payoff situations must preserve comparisons

[2] Such a smallest set always exists. To see this, let $\bar{\mathcal{E}}$ be the smallest superset of $\mathcal{E}$ satisfying symmetry and mixture properties; it exists and is the intersection of all such supersets, a family that is nonempty because it contains $\mathfrak{E}$. Now consider the family of comparison sets $\tilde{\mathcal{R}} \supseteq \mathcal{R}$ satisfying the transfer property with respect to $\bar{\mathcal{E}}$. This family is nonempty, since it contains $\mathfrak{R}$, and is closed under intersections. Hence its intersection is the smallest comparison set extending $\mathcal{R}$ through $\mathcal{E}$: any admissible auxiliary equivalence set contains $\bar{\mathcal{E}}$, so any comparison set extended through it also satisfies the transfer property with respect to $\bar{\mathcal{E}}$.

generated by all such convex combinations, and hence must be positive affine. This parallels the familiar uniqueness of expected-utility representations up to positive affine transformations (Von Neumann and Morgenstern, 2007; Herstein and Milnor, 1953).

Finally, we define what it means for agents to believe that others reason through a given set of equivalences. We do this using the knowledge sets introduced above:

Definition 5. *We say agent i believes that* ***agent j reasons according to the equivalence set*** $\mathcal{E} \subseteq \mathfrak{E}$ *if, for every $\omega \in \mathcal{K}_i$, the comparison set $\mathcal{R}_j^\omega$ is extended through $\mathcal{E}$. We say it is* ***common knowledge that agents reason according to*** $\mathcal{E}$ *if, for every $i, j \in \mathcal{I}$,*

$$\mathcal{K}_i \subseteq \left\{\omega \in \Omega : \mathcal{R}_j^\omega \text{ is extended through } \mathcal{E}\right\}.$$

Intuitively, this means agent i only considers possible awareness profiles in which agent j's comparison set contains all inferences that can be drawn from the equivalences in $\mathcal{E}$.

4 Strategic equivalence

We now apply our framework to strategically equivalent mechanisms. Strategic equivalence is a standard concept in the mechanism design literature: two mechanisms are strategically equivalent if one can be transformed into the other by relabeling each agent's actions, so that corresponding action profiles yield the same payoff consequences for all agents and types (see e.g. Vickrey (1961); Milgrom and Weber (1982); Krishna (2009)).[3] In our framework, this can be defined as follows.

Definition 6. *Mechanisms $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ are* ***strategically equivalent*** *if there exists a profile of bijections*

$$\alpha = (\alpha_i)_{i \in \mathcal{I}}, \qquad \alpha_i : \mathcal{A}_i \to \mathcal{A}_i'$$

such that, for every $i \in \mathcal{I}$, type $t_i \in \mathcal{T}_i$, and action profile $a \in \mathcal{A}$,

$$U_i^X[t_i, a] = U_i^{X'}[t_i, \alpha(a)].$$

A collection of mechanisms is a ***class of strategically equivalent mechanisms*** *if every pair of mechanisms in the collection is strategically equivalent.*

That is, relabeling each agent's action individually through α_i maps every action profile in X to one in X' that yields the same payoff for all agents and all types.[4] We illustrate this with

[3]Some notions of strategic equivalence allow one of the equivalent mechanisms to have payoff-equivalent duplicate actions, requiring bijections only after such duplicates are removed. We do not allow for this in order to simplify notation, but the definition could be straightforwardly amended to accommodate it.

[4]Note that strategic equivalence requires payoff equivalence, but does not require outcomes to coincide: two mechanisms may induce different outcomes as long as those differences are payoff-irrelevant.

the example of first-price and Dutch auctions whose strategic equivalence has been noted by Vickrey (1961).[5]

Example 3 (Vickrey (1961); Milgrom and Weber (1982))**.** *Consider an environment where a single good can be allocated to one of $N \geq 2$ agents, or remain unassigned. Each agent i has value $t_i \in \mathbb{R}$ for the good, and agents can be asked to make payments. This environment can be written as*

$$\left(\mathcal{I}, \mathcal{Y}, \mathbb{R}^{\mathcal{I}}, (u_i)_{i\in\mathcal{I}}\right), \text{ where } \mathcal{I} = \{1, \dots, N\}, \quad \mathcal{Y} := (\mathcal{I} \cup \{\varnothing\}) \times \mathbb{R}^{\mathcal{I}}_{+}, \quad u_i\big(t_i, (j, p)\big) = t_i \mathbb{1}\{j = i\} - p_i.$$

That is, an outcome specifies which agent receives the good and the payment made by each agent. The following two mechanisms for this environment are strategically equivalent.

First-price auction with bounded bids. *Each agent submits a bid in the interval $[0, \bar{b}]$. The good is allocated to one of the agents submitting the highest bid, according to some fixed tie-breaking rule, and the winner pays her bid. Formally,*

$$X^{\text{FPA}} = \left(\mathcal{A}^{\text{FPA}}, \Phi^{\text{FPA}}\right), \qquad \text{where} \qquad \mathcal{A}^{\text{FPA}} = [0, \bar{b}]^{\mathcal{I}}, \quad \Phi^{\text{FPA}}(a) = \big(w(a),\, a_{w(a)}\, e_{w(a)}\big),$$

where $w(a)$ denotes the agent selected as the winner from among $\arg\max_{j\in\mathcal{I}} a_j$ and $e_i \in \mathbb{R}^{\mathcal{I}}$ is the unit vector with a 1 in coordinate i and 0 elsewhere.

Dutch auction. *The price starts at $\bar{b}$ and falls linearly to 0 at speed $r > 0$. Each agent chooses a stopping time at which to claim the good. The good is allocated to one of the agents choosing the earliest stopping time, according to a fixed tie-breaking rule, and the winner pays the clock price at that time. Formally,*

$$X^{\text{Dutch}} = \left(\mathcal{A}^{\text{Dutch}}, \Phi^{\text{Dutch}}\right), \quad \text{where} \quad \mathcal{A}^{\text{Dutch}} = \left[0, \frac{\bar{b}}{r}\right]^{\mathcal{I}}, \quad \Phi^{\text{Dutch}}(a) = \big(\tilde{w}(a),\, (\bar{b} - r\, a_{\tilde{w}(a)})\, e_{\tilde{w}(a)}\big),$$

where $\tilde{w}(a)$ denotes the agent selected as the winner from among $\arg\min_{j\in\mathcal{I}} a_j$.

To understand this equivalence, note that by deciding when to stop the clock in the Dutch auction, a player is effectively choosing the price at which to claim the object. This is in turn equivalent to choosing a bid in a first-price auction. Indeed, these two decisions of player i can be connected with the following bijection:

$$\alpha_i^{\text{FPA}\to\text{Dutch}}(a_i) = \frac{\bar{b} - a_i}{r}.$$

This maps a bid a_i in the first-price auction to the moment when the Dutch clock reaches that same price. Therefore, for any bid profile in the first-price auction, the corresponding stopping-time profile in the Dutch auction yields the same winner and the same payment.

[5] The Dutch auction is dynamic; here we consider its reduced normal form where each bidder's strategy specifies the time at which to stop the clock.

Remark 4. *An analogous strategic equivalence holds between the sealed-bid second-price auction and the ascending auction, provided the latter does not reveal dropout information. If dropouts were observable, bidders could condition their play on others' dropout decisions, giving rise to strategies in the ascending auction that have no counterpart in the sealed-bid one. This issue does not arise for the first-price and Dutch auctions, since the Dutch auction ends the moment any bidder stops, so one cannot condition her play on any non-trivial information about others' actions.*

It is worth noting that bijections α characterize strategic equivalence entirely in terms of the mechanisms' *rules*: they establish that corresponding action profiles induce the same payoff consequences. Experimental evidence suggests that drawing such correspondences between mechanisms' rules can help agents strategize. Breitmoser and Schweighofer-Kodritsch (2022) find that performance in a sealed-bid second-price auction improves substantially when subjects are shown a descending-clock implementation that yields the same outcome. The authors argue that the clock format invites "continue versus stop" reasoning—an analogy to a setting that participants understand—which facilitates transferring those insights back to the sealed-bid format and reduces bidding errors.

We now discuss how communicating such rule-level correspondences to agents helps them transfer knowledge about equilibria from one mechanism to another. Suppose that two mechanisms are strategically equivalent and that all agents know an equilibrium of one of them, in the sense of Definition 2. If agents also understand that the two mechanisms are strategically equivalent, they should recognize that the equilibrium they know maps to an equilibrium of the other mechanism by appropriately relabeling actions. Consequently, the agents would then know the corresponding equilibrium of the other mechanism as well. This is captured by the following result.

Proposition 1. *Let $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ be strategically equivalent through α, and suppose it is common knowledge that agents reason according to the equivalence set*

$$\left\{(X, a, t_i) \underset{i}{\Leftrightarrow} (X', \alpha(a), t_i) : i \in \mathcal{I},\ t_i \in \mathcal{T}_i,\ a \in \mathcal{A}\right\}.$$

Then if there is an equilibrium of X for prior F that is commonly known, there is also an equilibrium of X' for prior F that is commonly known.

Note that the equivalence set in the proposition is always valid: since strategic equivalence requires $U_i^{X'}[t_i, \alpha(a)] = U_i^X[t_i, a]$, the condition of Theorem 1 is satisfied with the identity as the affine map. The substantive assumption is therefore not that these equivalences are correct, but that agents *recognize* them—that is, that the correspondence between the two mechanisms' rules has been explained to them, and that it is common knowledge that they reason accordingly.

The proof of the result is in the appendix. Intuitively, it shows that if σ is a commonly known equilibrium of mechanism X, its counterpart strategy profile where all the actions are appropriately relabeled, $\alpha_\# \sigma$, is a commonly known equilibrium of mechanism X'. This is because the equivalences generated by the action relabeling α, together with the reasoning closure in Definition 3, allow agents to transfer the payoff comparisons that verify equilibrium in the original mechanism to the relabeled strategy profile in the new one. Note also that this conclusion holds even though the equivalence set covers only pure action profiles, while the equilibrium strategies themselves may be mixed. This is because the mixture property in Definition 3 allows agents to extend pure-action equivalences to lotteries over actions.

Several features of this result deserve comment. First, what must be communicated to agents is the equivalence of the mechanisms' *rules*, and no information about equilibrium behavior needs to be conveyed. This distinction is practically relevant: while explaining to agents how they should play—which types should choose which actions and why—is a complex task, a correspondence between mechanisms' rules is a concrete and relatively simple object that could be communicated by a designer switching between mechanisms.

Second, the result is agnostic about how agents came to know the equilibrium of the original mechanism—the way they optimize their actions or form beliefs. It simply says that if agents have somehow acquired the relevant knowledge—formalized as the payoff comparisons in Definition 2—they do not need to repeat that process for the new mechanism.

Third, one may ask whether agents can also coordinate on the transferred equilibrium, particularly when multiple equilibria exist. While our framework does not model coordination directly, the communication of the rule-level correspondence between mechanisms provides a natural focal point: if agents understand that a certain strategy profile is the one they play in the original mechanism, and they understand how actions map across mechanisms, it is natural for them to expect that the corresponding profile will be played in the new one.

5 Strategic similarity

In the previous section we showed how understanding strategic equivalence allows agents to transfer knowledge across mechanisms. We now introduce a more permissive notion of *strategic similarity*, which generalizes strategic equivalence in two ways. First, while strategic equivalence only allows agents to relabel actions when moving from one mechanism to another, strategic similarity also allows them to reinterpret the meanings of *types*. Second, rather than requiring that utilities in the two mechanisms be equal after these remappings, it requires only that they be related through a valid equivalence set in the sense of Definition 4. This preserves a sense in which agents can transfer payoff comparisons across mechanisms while also letting us apply the notion in settings where payoffs lack an objective cardinal interpretation required to give meaning to cross-mechanism equalities.

Definition 7. *Let $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ be mechanisms for the same environment. The mechanisms X and X' are **strategically similar** if there exist profiles of bijections*

$$\alpha = (\alpha_i)_{i\in\mathcal{I}}, \qquad \alpha_i : \mathcal{A}_i \to \mathcal{A}'_i,$$

$$\tau = (\tau_i)_{i\in\mathcal{I}}, \qquad \tau_i : \mathcal{T}_i \to \mathcal{T}_i,$$

such that the equivalence set

$$\mathcal{E}^{X,X'}_{\alpha,\tau} := \left\{ (X, a, \tau_i(t_i)) \underset{i}{\Leftrightarrow} (X', \alpha(a), t_i) : i \in \mathcal{I},\ t_i \in \mathcal{T}_i,\ a \in \mathcal{A} \right\}$$

is valid. A collection of mechanisms is a ***class of strategically similar mechanisms*** *if every pair of mechanisms in the collection is strategically similar.*

What makes the definition restrictive is that all relabelings must be separable. Actions are relabeled agent by agent, as in strategic equivalence. Here the same is also true of types: each agent's type is remapped individually, and independently of the relabeling of actions.

By Theorem 1, this definition admits the following equivalent payoff representation.

Definition 7′. *Let $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ be mechanisms for the same environment. The mechanisms X and X' are* ***strategically similar*** *if there exist profiles of bijections*

$$\alpha = (\alpha_i)_{i\in\mathcal{I}}, \qquad \alpha_i : \mathcal{A}_i \to \mathcal{A}'_i,$$

$$\tau = (\tau_i)_{i\in\mathcal{I}}, \qquad \tau_i : \mathcal{T}_i \to \mathcal{T}_i,$$

and, for each $i \in \mathcal{I}$, functions

$$\kappa_i : \mathcal{T}_i \to \mathbb{R}_{++}, \qquad \lambda_i : \mathcal{T}_i \to \mathbb{R},$$

such that, for every $i \in \mathcal{I}$, every type $t_i \in \mathcal{T}_i$, and every pure action profile $a \in \mathcal{A}$,

$$U_i^{X'}[t_i, \alpha(a)] = \kappa_i(t_i) U_i^X[\tau_i(t_i), a] + \lambda_i(t_i).$$

This equivalent formulation is useful for two reasons. First, it provides a direct payoff condition that is easy to verify in applications. Second, it shows that strategic similarity is a self-contained notion of structural equivalence between mechanisms that can be used even in abstraction from our epistemic framework.

To build intuition for strategic similarity, we begin with a simple example.

Example 4. *Fix an environment with one agent who has value $t_1 \in \mathbb{R}$ for a good that may either be allocated to her or not, and where she may be charged a payment. This environment can be written as*

$$\left(\mathcal{I}, \mathcal{Y}, \mathbb{R}^{\mathcal{I}}, (u_i)_{i\in\mathcal{I}}\right), \quad \text{where} \qquad \mathcal{I} = \{1\}, \quad \mathcal{Y} := (\mathcal{I} \cup \{\varnothing\}) \times \mathbb{R}_+, \quad u_1\big(t_1, (y, p_1)\big) = t_1 \mathbb{1}\{y = 1\} - p_1.$$

Now consider a posted-price mechanism with price P in which the agent chooses whether to buy the good. If she buys it, she receives the good and pays P. Otherwise, she gets and pays nothing:

$$X_P = \left(\mathcal{A}^P, \Phi^P\right), \qquad \text{where} \quad \mathcal{A}^P = \mathcal{A}^P_1 = \{\textit{buy, do not buy}\}, \quad \Phi^P(a_1) = \begin{cases} (1, P) & \text{if } a_1 = \textit{buy}, \\ (\varnothing, 0) & \text{if } a_1 = \textit{do not buy}. \end{cases}$$

Then the family of posted-price mechanisms with prices $P \in \mathbb{R}_+$, $X_P = (\mathcal{A}^P, \Phi^P)$, forms a class of strategically similar mechanisms. Indeed, for any $P, P' \in \mathbb{R}_+$, let $\alpha_1^{P' \to P}$ be the identity map and choose

$$\tau_1^{P' \to P}(t_1) = t_1 + (P - P').$$

Then, for all types $t_1 \in \mathbb{R}$, we have

$$U_1^{X_{P'}}[t_1, buy] = t_1 - P' = \big(t_1 + (P - P')\big) - P = U_1^{X_P}[\tau_1^{P' \to P}(t_1), buy],$$

$$U_1^{X_{P'}}[t_1, do\ not\ buy] = 0 = U_1^{X_P}[\tau_1^{P' \to P}(t_1), do\ not\ buy].$$

Thus, the condition of Definition 7′ is satisfied with a scaling factor of 1 and an additive constant of 0. Intuitively, changing the posted price simply shifts the buyer's net value of buying from $t_1 - P'$ to $t_1 - P$. Thus, facing price P' as type t_1 is strategically the same as facing price P as type $t_1 + (P - P')$.

The example illustrates how the remappings in Definition 7 should be interpreted. The action map α identifies choices that play the same role in the two mechanisms, while the type map τ identifies agents whose payoff tradeoffs are the same after those choices are matched. In the posted-price example, the action correspondence is trivial, since "buy" and "do not buy" have the same meaning at every price, while the type remapping adjusts the buyer's value to account for the change in price.

Remark 5. *It is worth noting that an indirect mechanism, such as a first-price auction, is in general not strategically similar to the direct-revelation mechanism that implements the same outcome—that is, the mechanism in which each agent reports her type and the mechanism executes the equilibrium strategy on her behalf. Intuitively, this is because the "meaning" of an action in the indirect mechanism is fixed across all priors—a bid of b is a bid of b regardless of the distribution of opponents' types—while in the direct-revelation mechanism it is not: the action that a reported type maps to depends on the prior, since the equilibrium strategy the mechanism executes on the agent's behalf changes with beliefs about other participants. As a result, there is no single relabeling of actions that captures the correspondence between the two mechanisms independently of the prior, which is what strategic similarity requires.*

We now turn to discussing agents' understanding of strategic similarity. As in the previous section, we capture this by assuming it is common knowledge that they reason according to the equivalence set generated by the relevant correspondences between actions and types. As before, this equivalence set is automatically valid by Definition 7, so the substantive assumption is that agents recognize the correspondences—that is, that it has been explained to them how actions and types relate across the two mechanisms. We then get the following result:

Proposition 2. *Let $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$ be strategically similar through (α, τ), and suppose it is common knowledge that agents reason according to the equivalence set $\mathcal{E}_{\alpha,\tau}^{X,X'}$ from Definition 7. If for every prior $G \in \Delta(\mathcal{T})$ there exists a commonly known equilibrium of X under G, then for every prior $F \in \Delta(\mathcal{T})$, there also exists a commonly known equilibrium of X'.*

The proof is in the appendix. As in Proposition 1, the result shows that explaining the relevant equivalences to agents lets them transfer their equilibrium knowledge from one mechanism to the other. However, the result requires a stronger premise than the one for strategic equivalence: agents must know the equilibrium of the original mechanism X not just for a single prior, but for *every* prior over types. The reason is as follows. Under strategic similarity, agents can view mechanism X' as mechanism X with actions and types reinterpreted through the maps α, τ. Consequently, a prior F in X' corresponds to the *pushforward prior* $G = \tau_\# F$ in X, obtained by relabeling types according to τ. The payoff comparisons needed to certify the remapped strategy profile $\alpha_\# \sigma$ as an equilibrium of X' under F must therefore come from an equilibrium of X under the corresponding prior G. Nevertheless, if agents know an equilibrium for every prior of X, they are guaranteed to know one for this prior G, and can thus certify an equilibrium of X'.[6]

While requiring agents to know an equilibrium at every prior might seem strong, this condition has a natural interpretation: it says that agents understand some general method for finding equilibria for this mechanism given a belief over others' types. The result then says that this method also works in the other mechanism, once actions and types are appropriately relabeled.

More broadly, strategic similarity can also be understood in abstraction from our epistemic framework, as a way of certifying a common strategic structure across mechanisms. Indeed, such shared structure may itself help agents transfer knowledge across mechanisms, even when their reasoning differs somewhat from our model. This perspective will be useful when discussing applications: when strategic similarity holds, we explain intuitively how agents' problems relate between mechanisms; when it fails, we identify which strategic features are present in one mechanism but not in the other.

The following proposition illustrates the fact that strategic similarity preserves important strategic properties with the example of dominance solvability. The proof is in the appendix.

Definition 8. *A strategy $s_i : \mathcal{T}_i \to \Delta(\mathcal{A}_i)$ is **dominant** for agent i in mechanism X if*

$$s_i(t_i) \in \underset{A_i \in \Delta(\mathcal{A}_i)}{\arg\max}\, U_i^X[t_i, A_i \otimes \mu_{-i}] \qquad \textit{for every } t_i \in \mathcal{T}_i,\ \mu_{-i} \in \Delta(\mathcal{A}_{-i}).$$

Proposition 3. *Suppose X and X' are strategically similar. If mechanism X admits dominant strategies for all agents, then mechanism X' does too.*

5.1 Application: single-item auctions

We now apply the concept of strategic similarity to auction formats for the single-item environment from Example 3. We first characterize which standard auction families form strategically similar classes, and then turn to the implications for revenue-optimal auctions.

[6] Indeed, Proposition 2 could also be stated as saying that common knowledge of an equilibrium σ of X under G suffices for common knowledge of the equilibrium $\alpha_\# \sigma$ of X' under $\tau_\#^{-1} G$.

5.1.1 Reserve prices and entry costs. Fix $k \in \{1, \ldots, N\}$. For any profile of bids, let $a^{(k)}$ denote the kth-highest bid among participating bidders, with the convention that $a^{(k)} = 0$ if fewer than k bidders participate. Let $w(a)$ denote the bidder selected as the winner from among the highest participating bidders. Throughout this subsection, ties are broken by a fixed symmetric tie-breaking rule, so $\Phi(a)$ denotes the lottery over outcomes induced by this rule whenever the highest bid is tied.

(i) For $r \in \mathbb{R}_+$, the *kth-price auction with reserve r* is the mechanism $k\mathrm{PA}(r) = (\mathcal{A}^r, \Phi^r)$, where

$$\mathcal{A}_i^r = \{\varnothing\} \cup [r, \infty), \qquad \Phi^r(a) = \begin{cases} (\varnothing,\ \mathbf{0}) & \text{if } a_i = \varnothing \text{ for all } i, \\ \big(w(a),\ \max(a^{(k)}, r)\, e_{w(a)}\big) & \text{otherwise,} \end{cases}$$

where $w(a)$ and $a^{(k)}$ are computed over participating bidders (those with $a_i \neq \varnothing$), subject to the convention above. The ordinary *kth-price auction* is the special case $r = 0$.

(ii) For $c \in \mathbb{R}_+$, the *kth-price auction with entry cost c* is the mechanism $k\mathrm{EC}(c) = (\mathcal{A}^c, \Phi^c)$, where

$$\mathcal{A}_i^c = \{\varnothing\} \cup \mathbb{R}_+, \qquad \Phi^c(a) = \begin{cases} (\varnothing,\ \mathbf{0}) & \text{if } a_i = \varnothing \text{ for all } i, \\ \big(w(a),\ a^{(k)}\, e_{w(a)} + c \sum_i \mathbb{1}\{a_i \neq \varnothing\}\, e_i\big) & \text{otherwise,} \end{cases}$$

and $w(a)$ and $a^{(k)}$ are again computed over participating bidders. The winner pays her kth-price payment plus the entry cost, and every other participant pays only c.

We first observe that, for each fixed k, all kth-price auctions with reserves are closely related:

Theorem 2. *For each fixed k, the family of kth-price auctions with reserves, $\{k\mathrm{PA}(r) : r \in \mathbb{R}_+\}$, is a class of strategically similar mechanisms. Moreover, consider mechanisms $X = (\mathcal{A}, \Phi)$ with the following properties:*

1. ***Common action space.*** *There is a set $\mathcal{A}_0$ such that $\mathcal{A}_i = \mathcal{A}_0$ for every bidder i.*
2. ***Symmetry.*** *If a' is obtained from a by permuting bidders' actions under a, then the allocation probabilities and payments specified by $\Phi(a')$ are obtained from those specified by $\Phi(a)$ by the same permutation.*
3. ***Outside option.*** *There is an action $\varnothing \in \mathcal{A}_0$ such that, whenever bidder i chooses $\varnothing$, she never receives the object and never makes a payment. If at least one bidder chooses an action different from $\varnothing$, then the object is allocated with probability one among such bidders.*
4. ***No redundant actions.*** *For every bidder i and every distinct $a_i, a_i' \in \mathcal{A}_0$, there exists $a_{-i} \in \mathcal{A}_0^{N-1}$ such that $\Phi(a_i, a_{-i}) \neq \Phi(a_i', a_{-i})$.*

Then such a mechanism is strategically similar to a kth-price auction with some reserve if and only if it is strategically equivalent to a kth-price auction with possibly a different reserve.

Theorem 2 not only establishes that kth-price auctions with reserves form a strategically similar class, but also provides a sense in which this class is *maximal*. It says that if a designer wants to adjust a kth-price auction while preserving the aforementioned properties and strategic structure, the only parameter she can effectively vary is the reserve price. Indeed, the practice of fixing the auction format and varying only the reserve is widespread in industry. Sotheby's, Christie's, and eBay all commit to a standard auction format and use the reserve as their primary tuning parameter across sales.[7] In online advertising, major ad exchanges use the first-price auction as their standard format, with publishers optimizing revenue by adjusting reserve prices across impressions and keywords.[8]

While the proof is relegated to the appendix, we provide an intuition for why reserve prices preserve strategic similarity. Fix two reserves $r, r' \in \mathbb{R}_+$ and consider the bijections

$$\alpha_i^{r\to r'}(a) = \begin{cases} \varnothing & \text{if } a = \varnothing, \\ a + (r' - r) & \text{if } a \in [r, \infty), \end{cases} \qquad \tau_i^{r'\to r}(t_i) = t_i - (r' - r).$$

These maps shift every active bid up by $r' - r$ and every value down by the same amount. Since $\alpha^{r\to r'}$ preserves the set and ranking of active bids, it preserves the winner selected. Moreover, the kth-highest bid among participants rises by exactly $r' - r$, so the winner's payment $\max(a^{(k)}, r)$ also rises by $r' - r$. Thus, for every bidder i, type t_i, and action profile a:

$$\begin{aligned} U_i^{k\text{PA}(r')}[t_i, \alpha^{r\to r'}(a)] &= \mathbb{1}\{w(a) = i\}\left(t_i - (\max\{a^{(k)}, r\} + (r' - r))\right) \\ &= \mathbb{1}\{w(a) = i\}\left((t_i - (r' - r)) - \max\{a^{(k)}, r\}\right) = U_i^{k\text{PA}(r)}[\tau_i^{r'\to r}(t_i), a]. \end{aligned}$$

The condition of Definition 7′ is therefore satisfied with a scaling factor of 1 and an additive constant of 0. Note that this transformation is analogous to the case of posted-price mechanisms from Example 4. Intuitively, when the reserve rises from r to r', an agent can think of her choice in the new auction as the choice of a lower type in the old one.

Remark 6. *One might wonder whether the strategic similarity of auctions with reserves is special to the kth-price family. To answer this question, note first that the above construction of bijections α and τ works thanks to three properties of the kth-price format. First, only the winner pays; second, translating actions by the same constants preserves the winner; third, the winner's payment under reserve r, written as $\pi^r(a)$, satisfies:*

$$\pi^r(a) = \pi^0(a - r\mathbf{1}) + r.$$

[7] https://www.sothebys.com/en/buy-sell
https://www.christies.com/en/help/auction-help-library/how-do-auction-houses-work
https://www.ebay.com/help/selling/listings/selling-auctions/reserve-prices?id=4143.

[8] See e.g. https://support.google.com/admanager/answer/9298211.

As a result, changing the reserve can be absorbed by translating both bids and values. The action remapping $a_i \mapsto a_i - r$ preserves the allocation rule, while the type remapping $t_i \mapsto t_i - r$ offsets the payoff consequence of the additional r paid by the winner.

The same argument can therefore be applied to other formats with these properties, like the 1.5*-price auction, in which the winner pays an average of the first- and second-highest bids. The family of* 1.5*-price auctions with reserves is therefore also a class of strategically similar mechanisms.*

A version of this reasoning also extends to auctions where not only the winner pays. Consider, for instance, the all-pay auction, where every participant pays her own bid. There too shifting all bids up by $r' - r$ and all types down by the same amount preserves every bidder's payoff and yields a strategically similar family. However, the resulting family does not coincide with the family of all-pay auctions with reserves, understood as auctions with a minimum bid r: in an all-pay auction with reserve r, a loser who bids $b \geq r$ pays b, whereas the translated mechanism charges her $b + (r' - r)$.

The maximality statement in Theorem 2 also tells us which auctions are *not* strategically similar:

Corollary 1. *For $k \neq k'$, no kth-price auction with a reserve is strategically similar to any k'th-price auction with a reserve.*

For $k = 1$ and $k' = 2$ this also follows from Proposition 3, since the second-price auction is solvable in dominant strategies while the first-price auction is not; the theorem extends this observation to any $k \neq k'$. We also get the following result:

Corollary 2. *The family of kth-price auctions with entry costs, $\{k\mathrm{EC}(c) : c \in \mathbb{R}_+\}$, does not form a class of strategically similar mechanisms.*

To see why it holds, note that a kth-price auction with an entry cost of zero is also a kth price auction with a *reserve* of zero, and therefore Theorem 2 applies to it. Since kth-price auctions with entry costs $c > 0$ satisfy properties 1–4 but are not strategically equivalent to kth-price auctions with reserves, they cannot be strategically similar to the one with $c = 0$.

Consequently, kth-price auctions with reserves are more alike in their strategic structure than kth-price auctions with entry costs. This observation may not be obvious a priori. After all, both families contain the standard kth-price auction as a special case. It turns out, however, that introducing a strictly positive reserve to this baseline format changes the underlying strategic problem less than introducing a strictly positive entry cost. To get an intuition for this difference, consider how the agent's decision about whether to enter the auction changes as these parameters vary. First, note that the entry decision is just as simple with a reserve as it is without one: when the reserve is r, types $t_i < r$ should never enter while types $t_i > r$ always should. A strictly positive entry cost, by contrast, makes the participation decision strategically involved. Since an agent who enters and loses still pays c, entry is worthwhile only if the bidder expects a sufficiently high chance of winning, which in turn means that the entry decision depends on her beliefs about how likely other bidders are to enter and how aggressively they are likely to bid. In this sense, making the entry cost strictly positive introduces a qualitatively new strategic consideration, which breaks similarity to the baseline case.

Nevertheless, kth-price auctions with *strictly positive* entry costs are strategically similar to one another.

Proposition 4. *Fix $k \in \{1, \dots, N\}$. For any $c, c' \in \mathbb{R}_{++}$, the mechanisms $k\mathrm{EC}(c)$ and $k\mathrm{EC}(c')$ are strategically similar.*

Although belief-dependent entry arguably makes auctions with entry costs $c > 0$ harder than reserve-price auctions, this underscores that strategic similarity concerns the *structural similarity* of strategic problems, not their *difficulty*—strictly-positive-entry-cost auctions may be harder to play, but the reasoning needed to play in one transfers to any other.

5.1.2 Revenue-maximizing auctions. One might ask whether a revenue-maximizing designer who faces many possible bidder distributions can nevertheless restrict attention to a single strategically similar class of mechanisms. This is possible if attention is restricted to regular symmetric IPV priors: by Myerson (1981), the revenue-maximizing auction is then a common-reserve auction, and this family is strategically similar by Theorem 2.[9] However, once irregular symmetric priors are allowed, no single strategically similar class of symmetric mechanisms containing well-behaved auctions can achieve optimal revenue in general.

Theorem 3. *Let $\mathcal{C}$ be a strategically similar class of symmetric mechanisms. Suppose $\mathcal{C}$ contains a mechanism $X^0 = (\mathcal{A}^0, \Phi^0)$ with the following properties.*

1. ***Common bid space.*** *For every bidder i, $\mathcal{A}_i^0 = [\underline{a}, \infty) \cup \{\varnothing\}$ where $\varnothing$ is a nonparticipation action after which a bidder never receives the good and pays nothing.*
2. ***Highest bidder wins.*** *Whenever at least one bidder participates, the object is allocated among the highest participating bidders; ties are broken uniformly.*
3. ***Payment regularity.*** *For each bidder i and each action profile a, bidder i pays $\pi_w^0(a)$ in expectation conditional on winning and $\pi_\ell^0(a)$ in expectation conditional on losing. For every a_{-i}, $\pi_w^0(\cdot, a_{-i})$ is continuous and locally bounded on the set of bids with which bidder i can win, and $\pi_\ell^0(\cdot, a_{-i})$ is continuous and locally bounded on the set of bids with which bidder i can lose. Finally, at every highest-bid tie, the difference between bidder i's conditional winning and losing payments depends only on the common tied bid.*

Then $\mathcal{C}$ cannot implement a revenue-maximizing auction in symmetric equilibrium for every symmetric IPV prior.

[9] A prior F is regular if its virtual value $\varphi_F(v) = v - \frac{1-F(v)}{f(v)}$ is weakly increasing on its support.

To understand the result, recall that for irregular priors, the revenue-maximizing auction may allocate the object at random among a subset of bidders rather than always assigning it to the bidder with the highest value. This is the role of ironing: over certain intervals of values, the optimal auction treats distinct types as tied for allocation purposes. Thus, a strategically similar class that implements Myerson-optimal auctions for all symmetric IPV priors would have to contain mechanisms whose equilibria exhibit such pooling, as well as mechanisms whose equilibria assign the object to the highest bidder above a reserve. Theorem 3 shows that these two allocation structures are generally incompatible within any strategically similar class that contains a well-behaved highest-bid-wins auction.

Conditions on X^0 serve to rule out pathological constructions. Without some restriction of this kind, one could build mechanisms whose message spaces or payments are engineered to mimic arbitrary equilibrium allocations. Note also that the result does not require every mechanism in $\mathcal{C}$ to satisfy the properties above, but only requires that $\mathcal{C}$ contain one auction that does. In this sense, the tradeoff arises as soon as the designer sometimes wants to use a conventional auction format.

It is also worth noting that the positive result does not extend to asymmetric regular priors. When bidders draw values from different regular distributions $F_1, \ldots, F_N$, the Myerson-optimal auction can be implemented by a kth-price auction with bidder-specific reserves $r_i = \varphi_{F_i}^{-1}(0)$. However, the family of kth-price auctions with bidder-specific reserves does not in general form a strategically similar class. To see this, consider two bidders and two first-price auctions $\text{FPA}(r_1, r_2)$ and $\text{FPA}(r'_1, r'_2)$ with

$$r'_1 - r_1 \neq r'_2 - r_2.$$

The natural candidate bijections are the bidder-specific analogues of the translations used in Theorem 2:

$$\alpha_i(b_i) = b_i + (r'_i - r_i), \qquad \tau_i(t_i) = t_i - (r'_i - r_i).$$

These preserve each bidder's payoff conditional on winning, just as in the common-reserve case. However, when the shifts $r'_i - r_i$ differ across bidders, they do not preserve who actually wins. Indeed, this would require:

$$b_1 > b_2 \quad \Longleftrightarrow \quad b_1 + (r'_1 - r_1) > b_2 + (r'_2 - r_2)$$

for all feasible bid pairs, which fails whenever $r'_1 - r_1 \neq r'_2 - r_2$.

5.2 Application: procurement auctions

We now apply our framework to a procurement setting motivated by Allende et al. (2024). A buyer awards a contract to one of N sellers based on a score over price and quality. We consider two common price-scoring rules—a linear rule and a ratio rule—and show that changing the weight on price preserves strategic similarity under the linear rule but not the ratio rule.

Environment. A buyer awards a contract to one of N sellers. Each seller i has type $t_i = c_i \in \mathcal{T}_i = \mathbb{R}$ specifying her cost. An outcome $y = (w, p) \in \mathcal{Y} = \{1, \dots, N\} \times \mathbb{R}$ specifies the winning seller and the payment. Seller i's utility is

$$u_i\big(c_i, (w,p)\big) = \begin{cases} p - c_i & \text{if } w = i, \\ 0 & \text{if } w \neq i. \end{cases}$$

First-score auctions. In the mechanisms we study, each seller submits a price b_i. In addition to price, each seller i has a fixed quality score $q_i > 0$ which reflects observable characteristics such as past performance or delivery capacity. Given a weight $\lambda \in (0,1)$ and a price-scoring rule S_p, seller i's total score is

$$S_\lambda(a)_i = \lambda\, S_p(b_i, b_{-i}) + (1-\lambda)\, q_i.$$

The contract is awarded to one of the highest-scoring sellers, according to some fixed tie-breaking rule, and the winner is paid the price she submitted. Formally, letting $w(a)$ denote the seller selected as the winner from among the highest scorers, the outcome rule is

$$\Phi(a) = \big(w(a),\, b_{w(a)}\big).$$

To complete the specification of a first-score auction, it remains to choose the price-scoring rule S_p. We consider two such rules, each giving rise to a family of mechanisms indexed by the weight λ:

(i) The *linear-score* auction $L(\lambda)$, where $\mathcal{A}_i^{L(\lambda)} = \mathbb{R}$ and $S_\lambda(a)_i = -\lambda b_i + (1-\lambda)\, q_i$.

(ii) The *ratio-score* auction $R(\lambda)$, where $\mathcal{A}_i^{R(\lambda)} = \mathbb{R}_{++}$ and $S_\lambda(a)_i = \lambda \frac{\min_j b_j}{b_i} + (1-\lambda)\, q_i$.

The question is whether, within each family, varying the weight λ preserves strategic similarity. As it turns out, the linear family forms a class of strategically similar mechanisms, while the ratio family does not.

Proposition 5. *The linear family $\{L(\lambda) : \lambda \in (0,1)\}$ is a class of strategically similar mechanisms. However, unless the fixed quality scores q_i are all equal, the ratio family $\{R(\lambda) : \lambda \in (0,1)\}$ is not a class of strategically similar mechanisms.*

Intuitively, the linear family is strategically similar because changing λ can be absorbed by seller-specific translations of bids and costs. After the bid relabeling, all sellers' scores end up rescaled by the same positive constant, leaving the winner's identity unchanged. The corresponding change in the winner's payment is then exactly offset by the cost relabeling.

To see why the ratio family fails, notice that when λ is small, a sufficiently low-quality seller can never win: the price score is bounded above by λ, which in certain cases cannot outweigh other participants' quality bonuses. As λ increases, however, the set of bidders who have a chance of winning changes. But a mechanism in which some seller never wins cannot be strategically similar to one in which that seller wins for some bid profiles, since the affine payoff condition would force her payoff in the latter to be constant across all bid profiles, which it is not.

5.3 Application: input- vs. output-based pricing

In the last application, we consider a provider offering a service using limited processing capacity whose effectiveness varies over time. This setting captures a range of industries. A landowner leases acres whose yield depends on rainfall and soil conditions; a shipping line sells container slots whose value to shippers depends on routing and port congestion; a consulting firm sells associate time whose productivity depends on relevant institutional knowledge.

We compare two ways of pricing such services. Under *output pricing*, the provider charges based on the realized outcome, such as crop yield or the date at which cargo is delivered. Under *input pricing*, the provider charges for the resources the customer uses—per acre, per container slot, or per billable hour—with a tariff that does not vary with the current productivity of those resources. We show that capacity pricing, combined with proportional rationing of overdemand, generates a class of strategically similar mechanisms as the provider's efficacy varies. By contrast, output pricing generically does not.

Input-based pricing is common in settings of this kind. In community solar programs, for example, subscribers often pay for shares of installed capacity, even though the electricity generated by that capacity varies with environmental conditions.[10] Professional-services firms often charge per billable hour of staff time rather than based on concrete deliverables. A further example comes from U.S. agricultural land markets: the most common arrangement between a landowner and a tenant farmer is a *cash rent* lease, under which the tenant pays based on acreage regardless of how much the land produces in that season.[11] An alternative lease model, called a *crop share* lease, corresponds to output pricing. There, the landlord receives a fraction of the realized harvest, so payments vary directly with productivity.

Environment. A service provider has a resource capacity (normalized to 1) that she can allocate across $N \geq 2$ buyers. This capacity has an efficacy of $e \in \mathbb{R}_{++}$, that is, when capacity z_i is devoted to serving agent i, she receives the output $x_i = e\, z_i$. This setting is captured by the following environment:

$$(\mathcal{I}, \mathcal{Y}, \mathcal{T}, \mathcal{U}), \qquad \text{where} \qquad \mathcal{I} = \{1, \ldots, N\}, \qquad \mathcal{T}_i = \mathcal{V}, \qquad \mathcal{Y} = \mathbb{R}_+^{\mathcal{I}} \times \mathbb{R}^{\mathcal{I}},$$

and $\mathcal{V}$ is the set of all weakly increasing, weakly concave functions $v : \mathbb{R}_+ \to \mathbb{R}$ satisfying $v(0) = 0$. The agent's type is therefore a *function* specifying her value for different levels of delivered

[10] https://docs.nlr.gov/docs/fy23osti/86242.pdf

[11] https://www.canr.msu.edu/news/farm_land_rental_agreements_and_arrangements

output. An outcome $y = (x, p) \in \mathcal{Y}$ specifies the amount of output x_i allocated to each buyer and the payment p_i made by each buyer. Each buyer's Bernoulli utility function is

$$u_i(v_i, y) = v_i(x_i) - p_i.$$

We now introduce two families of mechanisms for this environment. Both will be defined in terms of a *regular price scheme*: a function $P : \mathbb{R}_+ \to \mathbb{R}_+$ that is continuous, strictly increasing, strictly convex, and satisfies $P(0) = 0$.

Input pricing mechanisms. Fix a regular price scheme P; a buyer who uses capacity z_i pays $P(z_i)$, regardless of the current efficacy. For each efficacy level $e > 0$, define the mechanism $X_e^{\text{in}} = ([0,1]^{\mathcal{I}}, \Phi_e^{\text{in}})$, where each buyer chooses a requested capacity $a_j \in [0,1]$. If orders exceed the provider's capacity, it is rationed proportionally; thus, given an action profile a, buyer i is allocated capacity

$$z_i(a) := a_i \min\left\{1, 1/\textstyle\sum_{j \in \mathcal{I}} a_j\right\},$$

and the outcome rule is

$$\Phi_e^{\text{in}}(a) = \big(x(a,e),\, p(a)\big), \qquad \text{where} \qquad x_i(a,e) = e\, z_i(a), \qquad p_i(a) = P\big(z_i(a)\big).$$

For every buyer i, type $v_i \in \mathcal{T}_i$, and action profile $a \in [0,1]^{\mathcal{I}}$,

$$U_i^{X_e^{\text{in}}}[v_i, a] = v_i\big(e\, z_i(a)\big) - P\big(z_i(a)\big).$$

Output pricing mechanisms. The tariff is defined over units of output rather than over capacity. Fix a regular price scheme P. For each efficacy level $e > 0$, define the mechanism $X_e^{\text{out}} = ([0,e]^{\mathcal{I}}, \Phi_e^{\text{out}})$, where each buyer requests output $a_i \in [0,e]$. Since capacity may still need to be rationed, given an action profile a, buyer i receives output

$$x_i(a,e) := a_i \min\left\{1, e/\textstyle\sum_{j \in \mathcal{I}} a_j\right\},$$

and the outcome rule is

$$\Phi_e^{\text{out}}(a) = \big(x(a,e),\, p(a,e)\big), \qquad \text{where} \qquad p_i(a,e) = P\big(x_i(a,e)\big).$$

Thus, for every buyer i, type $v_i \in \mathcal{T}_i$, and action profile $a \in \mathcal{A}_e^{\text{out}}$,

$$U_i^{X_e^{\text{out}}}[v_i, a] = v_i\big(x_i(a,e)\big) - P\big(x_i(a,e)\big).$$

Proposition 6. *For any regular price scheme P, the family of input-pricing mechanisms $\{X_e^{\text{in}} : e > 0\}$ is a class of strategically similar mechanisms. However, for a generic regular price scheme P, the family of output-pricing mechanisms $\{X_e^{\text{out}} : e > 0\}$ is not a class of strategically similar mechanisms.*

In fact, the conditions a price scheme must satisfy for the output-pricing family to be strategically similar are very stringent. In the proof we show that the family $\{X_e^{\text{out}} : e > 0\}$ can be a class of strategically similar mechanisms only if the price scheme is isoelastic, i.e. takes the form

$$P(x) = Ax^{\rho}$$

for some $A > 0$ and $\rho > 1$.

This suggests a sense in which input-based tariffs may be simpler for repeated customers. Over time, such customers may face different levels of need for the service (which in the model corresponds to different value functions over output) and thus know how to optimally choose orders under different valuations. Input-based tariffs make this knowledge easier to reuse when the effectiveness of the provider's capacity changes. A customer who understands how to choose capacity for different values of output can reinterpret a change in effectiveness as a change in her effective valuation, and then adjust her order accordingly.

6 Discussion

We develop a theory of how strategic reasoning transfers across mechanisms and formalize a sense in which genuinely different mechanisms can nevertheless be *strategically similar*. By characterizing classes of strategically similar mechanisms in various settings, we provide predictions that can be tested experimentally.

Our notion of similarity could be extended in several directions. One could, for instance, consider mechanisms that are "almost" strategically similar but differ in having additional actions that are dominated or redundant. Our current definition requires full bijections between action spaces, reflecting the baseline assumption that agents cannot easily identify such actions. One could relax this by first endowing agents with the ability to recognize dominated or redundant actions, having them eliminate these, and then requiring bijections only between the remaining ones. More substantially, one could extend the notion of similarity to compare mechanisms with different numbers of agents, or define one-directional reductions in which simpler mechanisms could be mapped to restricted versions of more complex ones. Such reductions would provide a natural notion of relative strategic complexity across mechanisms.

Finally, the framework we propose could be used to study strategic reasoning beyond the transfer of equilibrium knowledge. Our analysis endows agents with the ability to reason through payoff equivalences across mechanisms. One could, however, endow them with the ability to recognize structural properties of a *single* mechanism, such as symmetry or single-crossing conditions, and ask which strategic problems can be solved using those tools alone. This could provide a way of studying the difficulty of mechanisms outside dominant-strategy settings.

References

ALLENDE, C., J. P. ATAL, R. CARRIL, J. I. CUESTA, AND A. GONZÁLEZ-LIRA (2024): "Drivers of public procurement prices: Evidence from pharmaceutical markets," *International Journal of Industrial Organization*, 96, 103086, the 50th Annual Conference of European Association for Research in Industrial Economics, Rome, 2023.

AUMANN, R. AND A. BRANDENBURGER (1995): "Epistemic conditions for Nash equilibrium," *Econometrica: Journal of the Econometric Society*, 1161–1180.

AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4, 1236–1239.

BERNHEIM, B. D. (1984): "Rationalizable strategic behavior," *Econometrica: Journal of the Econometric Society*, 1007–1028.

BÖRGERS, T. AND J. LI (2019): "Strategically simple mechanisms," *Econometrica*, 87, 2003–2035.

BREITMOSER, Y. AND S. SCHWEIGHOFER-KODRITSCH (2022): "Obviousness around the clock," *Experimental Economics*, 25, 483–513.

BROOKS, B. AND S. DU (2025): "Simplicity and Portability in Mechanism Design: A Case for (and Against) the Worst Case," .

COOPER, D. J. AND J. H. KAGEL (2008): "Learning and transfer in signaling games," *Economic Theory*, 34, 415–439.

COPIC, J. AND A. GALEOTTI (2006): "Awareness as an equilibrium notion: Normal-form games," .

ELMES, S. AND P. J. RENY (1994): "On the strategic equivalence of extensive form games," *Journal of Economic Theory*, 62, 1–23.

FEINBERG, Y. (2021): "Games with unawareness," *The BE Journal of Theoretical Economics*, 21, 433–488.

GILBOA, I. AND D. SCHMEIDLER (1995): "Case-based decision theory," *The quarterly Journal of economics*, 110, 605–639.

HARSANYI, J. C. (1967): "Games with incomplete information played by "Bayesian" players, I–III Part I. The basic model," *Management science*, 14, 159–182.

HARSTAD, R. M. (2000): "Dominant strategy adoption and bidders' experience with pricing rules," *Experimental economics*, 3, 261–280.

HERSTEIN, I. N. AND J. MILNOR (1953): "An axiomatic approach to measurable utility," *Econometrica, Journal of the Econometric Society*, 291–297.

JEHIEL, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81–104.

KRISHNA, V. (2009): *Auction theory*, Academic press.

LI, S. (2017): "Obviously strategy-proof mechanisms," *American Economic Review*, 107, 3257–3287.

——— (2024): "Designing Simple Mechanisms," *Journal of Economic Perspectives*, 38, 175–92.

MCKINSEY, J. (2016): "Isomorphism of games and strategic equivalence," *Contributions to the Theory of Games*, 1, 117–130.

MERTENS, J.-F. AND S. ZAMIR (1985): "Formulation of Bayesian analysis for games with incomplete information," *International journal of game theory*, 14, 1–29.

MILGROM, P. R. AND R. J. WEBER (1982): "A theory of auctions and competitive bidding," *Econometrica: Journal of the Econometric Society*, 1089–1122.

MORRIS, S. AND T. UI (2005): "Generalized potentials and robust sets of equilibria," *Journal of Economic Theory*, 124, 45–78.

MOULIN, H. AND J.-P. VIAL (1978): "Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon," *International Journal of Game Theory*, 7, 201–221.

MYERSON, R. B. (1981): "Optimal auction design," *Mathematics of operations research*, 6, 58–73.

PEARCE, D. G. (1984): "Rationalizable strategic behavior and the problem of perfection," *Econometrica: Journal of the Econometric Society*, 1029–1050.

PYCIA, M. AND P. TROYAN (2023): "A theory of simplicity in games and mechanism design," *Econometrica*, 91, 1495–1526.

RICK, S. AND R. A. WEBER (2010): "Meaningful learning and transfer of learning in games played repeatedly without feedback," *Games and Economic Behavior*, 68, 716–730.

SAMUELSON, L. (2001): "Analogies, adaptation, and anomalies," *Journal of Economic Theory*, 97, 320–366.

TEWOLDE, E. AND V. CONITZER (2021): "Game transformations that preserve nash equilibria or best-response sets," *arXiv preprint arXiv:2111.00076*.

VICKREY, W. (1961): "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, 16, 8–37.

Von Neumann, J. and O. Morgenstern (2007): "Theory of games and economic behavior: 60th anniversary commemorative edition," in *Theory of games and economic behavior*, Princeton university press.

A Omitted proofs

A.1 Proof of Theorem 1

($\Leftarrow$). Suppose the affine condition holds and fix any $\mathcal{R} \subseteq \mathfrak{R}$ containing only true comparisons. We show that the smallest $\hat{\mathcal{R}} \supseteq \mathcal{R}$ that is extended through $\mathcal{E}$ contains only true comparisons.

Let $\mathfrak{R}^{\text{true}} \subseteq \mathfrak{R}$ denote the set of all true payoff comparisons and $\mathcal{E}^*$ be the set of all payoff equivalences $(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t'_j)$ such that

$$U_j^{X'}[t'_j, A'] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A]\big).$$

By assumption, $\mathcal{E} \subseteq \mathcal{E}^*$. We show that the pair $(\mathfrak{R}^{\text{true}}, \mathcal{E}^*)$ satisfies the transfer, symmetry, and mixture properties in Definition 3. First, consider the transfer property. Suppose

$$\big\{(X, A^1, t_j) \underset{j}{\Leftrightarrow} (X', A'^1, t'_j),\ (X, A^2, t_j) \underset{j}{\Leftrightarrow} (X', A'^2, t'_j)\big\} \subseteq \mathcal{E}^*, \quad \big\{U_j^X[t_j, A^1] \geq U_j^X[t_j, A^2]\big\} \in \mathfrak{R}^{\text{true}}.$$

Then

$$U_j^{X'}[t'_j, A'^1] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A^1]\big), \qquad U_j^{X'}[t'_j, A'^2] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A^2]\big).$$

Since the affine transformation is strictly increasing, it follows that $U_j^{X'}[t'_j, A'^1] \geq U_j^{X'}[t'_j, A'^2]$, so the transfer produces only true comparisons. Second, note $\mathcal{E}^*$ is closed under symmetry because, by (1),

$$U_j^{X'}[t'_j, A'] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A]\big) \quad \Leftrightarrow \quad U_j^X[t_j, A] = \ell^j_{(X',t'_j)\to(X,t_j)}\big(U_j^{X'}[t'_j, A']\big).$$

Third, consider mixtures. Let (S, μ) be a probability space and suppose that, for every $s \in S$,

$$\big\{(X, A^s, t_j) \underset{j}{\Leftrightarrow} (X', A'^s, t'_j)\big\} \in \mathcal{E}^*.$$

Then, for every $s \in S$,

$$U_j^{X'}[t'_j, A'^s] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A^s]\big).$$

Since expected utility is affine in the mixed action profile and $\ell^j_{(X,t_j)\to(X',t'_j)}$ is affine,

$$\begin{aligned} U_j^{X'}\left[t'_j, \int_S A'^s \, d\mu\right] &= \int_S U_j^{X'}[t'_j, A'^s]\, d\mu \\ &= \int_S \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A^s]\big)\, d\mu \\ &= \ell^j_{(X,t_j)\to(X',t'_j)}\left(\int_S U_j^X[t_j, A^s]\, d\mu\right) = \ell^j_{(X,t_j)\to(X',t'_j)}\left(U_j^X\left[t_j, \int_S A^s\, d\mu\right]\right). \end{aligned}$$

Therefore the mixed equivalence also belongs to $\mathcal{E}^*$.

Thus $\mathfrak{R}^{\text{true}}$ is extended through $\mathcal{E}$, using $\mathcal{E}^*$ as the associated equivalence set. Since $\mathcal{R} \subseteq \mathfrak{R}^{\text{true}}$, the smallest $\hat{\mathcal{R}} \supseteq \mathcal{R}$ that is extended through $\mathcal{E}$ is contained in $\mathfrak{R}^{\text{true}}$. Hence it contains only true comparisons, so $\mathcal{E}$ is valid.

$(\Rightarrow)$. Suppose $\mathcal{E}$ is valid and let $\overline{\mathcal{E}}$ be the smallest set of equivalences containing $\mathcal{E}$ and closed under the symmetry and mixture properties. Fix an agent j, mechanisms $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$, and types $t_j, t'_j \in \mathcal{T}_j$. Define

$$S := \left\{\left(U_j^X[t_j, A], U_j^{X'}[t'_j, A']\right) : \{(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t'_j)\} \in \overline{\mathcal{E}}\right\} \subseteq \mathbb{R}^2.$$

Step 1. Validity implies order preservation. Take any $(x_1, y_1), (x_2, y_2) \in S$ and let their corresponding equivalences in $\overline{\mathcal{E}}$ be

$$(X, A^1, t_j) \underset{j}{\Leftrightarrow} (X', A'^1, t'_j), \qquad (X, A^2, t_j) \underset{j}{\Leftrightarrow} (X', A'^2, t'_j),$$

that is,

$$x_1 = U_j^X[t_j, A^1], \quad x_2 = U_j^X[t_j, A^2], \quad y_1 = U_j^{X'}[t'_j, A'^1], \quad y_2 = U_j^{X'}[t'_j, A'^2].$$

We claim that $x_1 \geq x_2$ if and only if $y_1 \geq y_2$. Suppose first that $x_1 \geq x_2$. Then $U_j^X[t_j, A^1] \geq U_j^X[t_j, A^2]$ is a true comparison. Let $\mathcal{R}$ be the singleton set containing this comparison. Now let $\tilde{\mathcal{R}} \supseteq \mathcal{R}$ be any comparison set extended through $\mathcal{E}$, with associated equivalence set $\tilde{\mathcal{E}} \supseteq \mathcal{E}$. Since $\tilde{\mathcal{E}}$ is closed under symmetry and mixtures, we have $\overline{\mathcal{E}} \subseteq \tilde{\mathcal{E}}$. Therefore the two equivalences above belong to $\tilde{\mathcal{E}}$. By the transfer property,

$$\left\{U_j^{X'}[t'_j, A'^1] \geq U_j^{X'}[t'_j, A'^2]\right\} \in \tilde{\mathcal{R}}.$$

Since this holds for every such extended $\tilde{\mathcal{R}}$, this comparison belongs to the smallest $\hat{\mathcal{R}} \supseteq \mathcal{R}$ that is extended through $\mathcal{E}$. By validity of $\mathcal{E}$, this comparison is true. Hence $y_1 \geq y_2$.

Conversely, suppose $y_1 \geq y_2$. Then $U_j^{X'}[t'_j, A'^1] \geq U_j^{X'}[t'_j, A'^2]$ is a true comparison. An analogous argument then holds because the reverse equivalences

$$(X', A'^1, t'_j) \underset{j}{\Leftrightarrow} (X, A^1, t_j), \qquad (X', A'^2, t'_j) \underset{j}{\Leftrightarrow} (X, A^2, t_j)$$

belong to $\overline{\mathcal{E}}$ by the symmetry property.

Step 2. Order preservation makes S the graph of a function. Observe that if two points in S have the same first coordinate, then they have the same second coordinate: if $x_1 = x_2$, then both $x_1 \geq x_2$ and $x_2 \geq x_1$, so Step 1 gives $y_1 \geq y_2$ and $y_2 \geq y_1$. Thus $y_1 = y_2$. Hence there is a set $D \subseteq \mathbb{R}$ and a function $f : D \to \mathbb{R}$ such that

$$S = \{(x, f(x)) : x \in D\}.$$

Moreover, Step 1 implies that f is strictly increasing whenever D contains more than one point.

Step 3. Mixtures force the induced function to be affine. If D is empty, choose

$$\ell^j_{(X,t_j)\to(X',t'_j)}(z) := z.$$

If D is a singleton, say $D = \{x_0\}$ and $f(x_0) = y_0$, choose any $\kappa > 0$ and set

$$\ell^j_{(X,t_j)\to(X',t'_j)}(z) := \kappa z + (y_0 - \kappa x_0) \quad \Longrightarrow \quad \ell^j_{(X,t_j)\to(X',t'_j)}(x_0) = y_0.$$

Now suppose D contains at least two points. Take distinct $x_0, x_1 \in D$, and choose equivalences in $\overline{\mathcal{E}}$ corresponding to $(x_0, f(x_0))$ and $(x_1, f(x_1))$. For any $\alpha \in [0,1]$, the mixture property gives

$$\big(\alpha x_0 + (1-\alpha)x_1,\ \alpha f(x_0) + (1-\alpha) f(x_1)\big) \in S.$$

Therefore

$$f\big(\alpha x_0 + (1-\alpha)x_1\big) = \alpha f(x_0) + (1-\alpha) f(x_1).$$

Assume without loss of generality that $x_0 < x_1$, and define

$$\kappa := \frac{f(x_1) - f(x_0)}{x_1 - x_0}, \qquad \lambda := f(x_0) - \kappa x_0.$$

Since f is strictly increasing, $\kappa > 0$. We claim that $f(x) = \kappa x + \lambda$ for every $x \in D$.

If $x \in [x_0, x_1] \cap D$, then $x = \alpha x_1 + (1-\alpha)x_0$ for some $\alpha \in [0,1]$, so

$$f(x) = \alpha f(x_1) + (1-\alpha) f(x_0) = \kappa x + \lambda.$$

If $x > x_1$, then $x_1 = \alpha x + (1-\alpha)x_0$ for some $\alpha \in (0,1)$. Hence

$$f(x_1) = \alpha f(x) + (1-\alpha) f(x_0),$$

which implies $f(x) = \kappa x + \lambda$. The case $x < x_0$ is identical, using that x_0 is a convex combination of x and x_1. Thus f agrees on all of D with the positive affine function

$$\ell^j_{(X,t_j)\to(X',t'_j)}(z) := \kappa z + \lambda.$$

Step 4. The affine map represents every original equivalence. Since $\mathcal{E} \subseteq \overline{\mathcal{E}}$, every original equivalence

$$(X, A, t_j) \underset{j}{\Leftrightarrow} (X', A', t'_j)$$

in $\mathcal{E}$ satisfies

$$U_j^{X'}[t'_j, A'] = \ell^j_{(X,t_j)\to(X',t'_j)}\big(U_j^X[t_j, A]\big).$$

Step 5. The reverse affine map can be chosen as the inverse. Finally, choose the affine map for the reverse ordered pair to be the inverse:

$$\ell^j_{(X',t'_j)\to(X,t_j)} = \left(\ell^j_{(X,t_j)\to(X',t'_j)}\right)^{-1}.$$

This is valid because the symmetry part of Definition 3 puts the reverse equivalences in $\overline{\mathcal{E}}$, and the order-preservation argument above identifies exactly the inverse payoff relation. In the self-pair case $(X, t_j) = (X', t'_j)$, the same argument applied to an equivalence and its symmetric counterpart implies that every self-equivalence in $\overline{\mathcal{E}}$ preserves payoff exactly; we may therefore choose the identity map, which is its own inverse.

A.2 Proof of Proposition 1

Let σ be an equilibrium of X for prior F that is commonly known. For each mixed action $A_i \in \Delta(\mathcal{A}_i)$, write $(\alpha_i)_\# A_i$ for its push-forward under α_i. Define a strategy profile σ' in X' by

$$\sigma'_i(t_i) := (\alpha_i)_\# \sigma_i(t_i) \qquad \text{for every } i \text{ and } t_i \in \mathcal{T}_i.$$

We show that it is commonly known that σ' is an equilibrium of X' for prior F.

First note that, since X and X' are strategically equivalent through α, for every agent j, type t_j, and mixed action profile $A \in \Delta(\mathcal{A})$,

$$U_j^{X'}[t_j, \alpha_\# A] = U_j^X[t_j, A].$$

Hence σ' is an equilibrium of X' for prior F: every deviation in X' is the α_j-push-forward of a deviation in X, and the relevant expected payoffs coincide.

It remains to show that this equilibrium is commonly known. Let

$$\mathcal{E}^{\alpha} := \big\{(X, a, t_i) \underset{i}{\Leftrightarrow} (X', \alpha(a), t_i) : i \in \mathcal{I},\ t_i \in \mathcal{T}_i,\ a \in \mathcal{A}\big\}.$$

Fix agents i, h, j, a type $t_j \in \mathcal{T}_j$, a mixed deviation $A'_j \in \Delta(\mathcal{A}'_j)$, and an awareness profile $\omega \in \mathcal{K}_i$. Since α_j is a bijection, there exists $A_j \in \Delta(\mathcal{A}_j)$ such that $A'_j = (\alpha_j)_\# A_j$. By construction of σ',

$$A_{-j}(\sigma', F \mid t_j) = (\alpha_{-j})_\# A_{-j}(\sigma, F \mid t_j).$$

Since it is commonly known that σ is an equilibrium of X for prior F, we have

$$\Big\{U_j^X\big[t_j, \sigma_j(t_j) \otimes A_{-j}(\sigma, F \mid t_j)\big] \geq U_j^X\big[t_j, A_j \otimes A_{-j}(\sigma, F \mid t_j)\big]\Big\} \in \mathcal{R}_h^{\omega}.$$

Since it is common knowledge that agents reason according to $\mathcal{E}^{\alpha}$, the set $\mathcal{R}_h^{\omega}$ is extended through $\mathcal{E}^{\alpha}$. Let $\hat{\mathcal{E}} \supseteq \mathcal{E}^{\alpha}$ witness this extension. The pure-profile equivalences in $\mathcal{E}^{\alpha}$, together with the mixture property in Definition 3, imply that

$$\big(X, \sigma_j(t_j) \otimes A_{-j}(\sigma, F \mid t_j), t_j\big) \underset{j}{\Leftrightarrow} \big(X', \sigma'_j(t_j) \otimes A_{-j}(\sigma', F \mid t_j), t_j\big)$$

and

$$\big(X, A_j \otimes A_{-j}(\sigma, F \mid t_j), t_j\big) \underset{j}{\Leftrightarrow} \big(X', A'_j \otimes A_{-j}(\sigma', F \mid t_j), t_j\big)$$

belong to $\hat{\mathcal{E}}$. Applying the transfer property in Definition 3 therefore gives

$$\Big\{U_j^{X'}\big[t_j, \sigma'_j(t_j) \otimes A_{-j}(\sigma', F \mid t_j)\big] \geq U_j^{X'}\big[t_j, A'_j \otimes A_{-j}(\sigma', F \mid t_j)\big]\Big\} \in \mathcal{R}_h^{\omega}.$$

A.3 Proof of Proposition 2

Fix $F \in \Delta(\mathcal{T})$. Since $\tau_\# F \in \Delta(\mathcal{T})$, by hypothesis there exists an equilibrium σ of X for prior $\tau_\# F$ that is commonly known. For each mixed action $A_i \in \Delta(\mathcal{A}_i)$, write $(\alpha_i)_\# A_i$ for its push-forward under α_i. Define a strategy profile σ' in X' by

$$\sigma'_i(t_i) := (\alpha_i)_\# \sigma_i(\tau_i(t_i)) \qquad \text{for every } i \text{ and } t_i \in \mathcal{T}_i.$$

We show that it is commonly known that σ' is an equilibrium of X' for prior F.

First, σ' is an equilibrium of X' for prior F. Indeed, by strategic similarity and Theorem 1, for each agent j and type t_j, the payoff in X' from the α-image of any mixed profile is a positive affine transformation of the corresponding payoff in X for type $\tau_j(t_j)$. Thus, the set of best responses is preserved under α_j. Since σ is an equilibrium of X for prior $\tau_\# F$, the profile σ' is an equilibrium of X' for prior F.

It remains to show that this equilibrium is commonly known. Fix agents i, h, j, a type $t_j \in \mathcal{T}_j$, a mixed deviation $A'_j \in \Delta(\mathcal{A}'_j)$, and an awareness profile $\omega \in \mathcal{K}_i$. Since α_j is a bijection, there exists $A_j \in \Delta(\mathcal{A}_j)$ such that $A'_j = (\alpha_j)_\# A_j$. Moreover, by construction of σ' and by the definition of the push-forward prior $\tau_\# F$,

$$A_{-j}(\sigma', F \mid t_j) = (\alpha_{-j})_\# A_{-j}(\sigma, \tau_\# F \mid \tau_j(t_j)).$$

Since it is commonly known that σ is an equilibrium of X for prior $\tau_\# F$, we have

$$\Big\{ U_j^X\big[\tau_j(t_j), \sigma_j(\tau_j(t_j)) \otimes A_{-j}(\sigma, \tau_\# F \mid \tau_j(t_j))\big] \geq U_j^X\big[\tau_j(t_j), A_j \otimes A_{-j}(\sigma, \tau_\# F \mid \tau_j(t_j))\big] \Big\} \in \mathcal{R}_h^\omega.$$

Since it is common knowledge that agents reason according to $\mathcal{E}_{\alpha,\tau}^{X,X'}$, the set $\mathcal{R}_h^\omega$ is extended through $\mathcal{E}_{\alpha,\tau}^{X,X'}$. Let $\hat{\mathcal{E}} \supseteq \mathcal{E}_{\alpha,\tau}^{X,X'}$ witness this extension. The pure-profile equivalences in $\mathcal{E}_{\alpha,\tau}^{X,X'}$, together with the mixture property in Definition 3, imply that

$$\big(X, \sigma_j(\tau_j(t_j)) \otimes A_{-j}(\sigma, \tau_\# F \mid \tau_j(t_j)), \tau_j(t_j)\big) \underset{j}{\Leftrightarrow} \big(X', \sigma'_j(t_j) \otimes A_{-j}(\sigma', F \mid t_j), t_j\big)$$

and

$$\big(X, A_j \otimes A_{-j}(\sigma, \tau_\# F \mid \tau_j(t_j)), \tau_j(t_j)\big) \underset{j}{\Leftrightarrow} \big(X', A'_j \otimes A_{-j}(\sigma', F \mid t_j), t_j\big)$$

belong to $\hat{\mathcal{E}}$. Applying the transfer property in Definition 3 therefore gives

$$\Big\{ U_j^{X'}\big[t_j, \sigma'_j(t_j) \otimes A_{-j}(\sigma', F \mid t_j)\big] \geq U_j^{X'}\big[t_j, A'_j \otimes A_{-j}(\sigma', F \mid t_j)\big] \Big\} \in \mathcal{R}_h^\omega.$$

A.4 Proof of Proposition 3

Let $X = (\mathcal{A}, \Phi)$ and $X' = (\mathcal{A}', \Phi')$, and suppose that X admits dominant strategies for all agents. Let α, τ and κ, λ be the bijections and functions from the equivalent definition of strategic similarity. Then, for every agent i, every type $t_i \in \mathcal{T}_i$, and every pure action profile $a \in \mathcal{A}$,

$$U_i^{X'}[t_i, \alpha(a)] = \kappa_i(t_i) U_i^X[\tau_i(t_i), a] + \lambda_i(t_i).$$

Since X admits dominant strategies for all agents, for each agent i choose a dominant strategy $s_i : \mathcal{T}_i \to \Delta(\mathcal{A}_i)$, and define $s'_i : \mathcal{T}_i \to \Delta(\mathcal{A}'_i)$ by $s'_i(t_i) := (\alpha_i)_\# s_i\big(\tau_i(t_i)\big)$. We show that s'_i is dominant for agent i in mechanism X'. Fix $t_i \in \mathcal{T}_i$ and $\mu'_{-i} \in \Delta(\mathcal{A}'_{-i})$, and let $\mu_{-i} := (\alpha_{-i}^{-1})_\# \mu'_{-i}$. By strategic similarity and linearity, for every $A_i \in \Delta(\mathcal{A}_i)$,

$$U_i^{X'}\big[t_i, (\alpha_i)_\# A_i \otimes \mu'_{-i}\big] = \kappa_i(t_i) U_i^X\big[\tau_i(t_i), A_i \otimes \mu_{-i}\big] + \lambda_i(t_i).$$

Because s_i is dominant for agent i in mechanism X, for every $A_i \in \Delta(\mathcal{A}_i)$,

$$U_i^X\big[\tau_i(t_i), s_i(\tau_i(t_i)) \otimes \mu_{-i}\big] \geq U_i^X\big[\tau_i(t_i), A_i \otimes \mu_{-i}\big].$$

Since $\kappa_i(t_i) > 0$, multiplying by $\kappa_i(t_i)$ and adding $\lambda_i(t_i)$ preserves the inequality. Translating back to mechanism X', this gives

$$U_i^{X'}\big[t_i, s_i'(t_i) \otimes \mu'_{-i}\big] \geq U_i^{X'}\big[t_i, (\alpha_i)_\# A_i \otimes \mu'_{-i}\big]$$

for every $A_i \in \Delta(\mathcal{A}_i)$. Now, fix any $A_i' \in \Delta(\mathcal{A}_i')$ and let $A_i := (\alpha_i^{-1})_\# A_i'$. Then $(\alpha_i)_\# A_i = A_i'$, so

$$U_i^{X'}\big[t_i, s_i'(t_i) \otimes \mu'_{-i}\big] \geq U_i^{X'}\big[t_i, A_i' \otimes \mu'_{-i}\big].$$

Since t_i, μ'_{-i}, and A_i' were arbitrary, s_i' is dominant for agent i in mechanism X'.

A.5 Proof of Theorem 2

We first show that the family $\{k\text{PA}(r) : r \in \mathbb{R}_+\}$ is a class of strategically similar mechanisms. Fix $k \leq N$ and $r, r' \in \mathbb{R}_+$. Define

$$\alpha_i(\varnothing) = \varnothing, \quad \alpha_i(b) = b + (r' - r) \text{ for } b \in [r, \infty), \quad \text{and} \quad \tau_i(t) = t - (r' - r), \quad \kappa_i(t) = 1, \quad \lambda_i(t) = 0.$$

For every bid profile a, the map α preserves the set of participants, the set of highest bidders, and hence the distribution over winners induced by the symmetric tie-breaking rule. Moreover, whenever the object is allocated, $\pi^{r'}(\alpha(a)) = \pi^r(a) + (r' - r)$. Therefore, for every bidder i, type t_i, and action profile a,

$$U_i^{k\text{PA}(r')}[t_i, \alpha(a)] = U_i^{k\text{PA}(r)}[t_i - (r' - r), a].$$

Thus, by Definition 7′, $k\text{PA}(r)$ and $k\text{PA}(r')$ are strategically similar.

We now show the maximality claim. Let $X = (\mathcal{A}, \Phi)$ satisfy the four properties, and suppose that X is strategically similar to $k\text{PA}(r)$. By the affine formulation of strategic similarity, there exist bijections

$$\alpha_i : \mathcal{A}_i^r \to \mathcal{A}_0, \qquad \tau_i : \mathbb{R} \to \mathbb{R},$$

and functions

$$\kappa_i : \mathbb{R} \to \mathbb{R}_{++}, \qquad \lambda_i : \mathbb{R} \to \mathbb{R},$$

such that, for every i, every $t \in \mathbb{R}$, and every $a \in \mathcal{A}^r$,

$$U_i^X[t, \alpha(a)] = \kappa_i(t) U_i^{k\text{PA}(r)}[\tau_i(t), a] + \lambda_i(t). \tag{2}$$

Step 1: $\alpha_i(\varnothing) = \varnothing$ and $\lambda_i \equiv 0$. Fix i. If $a_i = \varnothing$, then $U_i^{k\mathrm{PA}(r)}[\tau_i(t), a] = 0$, so by (2),

$$U_i^X[t, \alpha_i(\varnothing), \alpha_{-i}(a_{-i})] = \lambda_i(t) \quad \text{for all } t, a_{-i}. \tag{3}$$

Because each α_j is onto, this implies that

$$U_i^X[t, \alpha_i(\varnothing), b_{-i}] = \lambda_i(t) \quad \text{for all } t, b_{-i} \in \mathcal{A}_0^{N-1}. \tag{4}$$

If $\alpha_i(\varnothing) \neq \varnothing$, then when bidder i chooses $\alpha_i(\varnothing)$ and all other bidders choose the true outside action $\varnothing$, the outside-option property implies that bidder i receives the object with probability one. Hence the left side of (4) has slope one in t. But when all bidders choose $\alpha_i(\varnothing)$, symmetry implies that bidder i's allocation probability is $1/N$, so the left side of (4) has slope $1/N$ in t. Since $N \geq 2$, this contradicts (4). Therefore $\alpha_i(\varnothing) = \varnothing$. Substituting back into (4) and using the outside-option property of X gives:

$$\lambda_i(t) = 0 \quad \text{for all } t. \tag{5}$$

Step 2: Zero allocation in kPA(r) implies zero allocation and payment in X. Throughout, we write $q_i^Y(a)$ and $p_i^Y(a)$ for bidder i's interim allocation probability and interim expected payment under action profile a in mechanism Y. Fix a bidder i. If $q_i^{k\mathrm{PA}(r)}(a) = 0$, then $U_i^{k\mathrm{PA}(r)}[\tau_i(t), a] = 0$ for every t. Hence, by (2) and (5), $U_i^X[t, \alpha(a)] = 0$ for all t. Since $U_i^X[t, \alpha(a)] = tq_i^X(\alpha(a)) - p_i^X(\alpha(a))$, we therefore have

$$q_i^{k\mathrm{PA}(r)}(a) = 0 \quad \Longrightarrow \quad q_i^X(\alpha(a)) = p_i^X(\alpha(a)) = 0. \tag{6}$$

Step 3: The scaling function κ_i is constant and τ_i is affine. Take two profiles $a, \tilde{a}$ at which bidder i strictly wins in kPA(r), with different payments $\pi^r(a) \neq \pi^r(\tilde{a})$. From (2) and (5), we get:

$$tq_i^X(\alpha(a)) - p_i^X(\alpha(a)) = \kappa_i(t)\big(\tau_i(t) - \pi^r(a)\big), \quad tq_i^X(\alpha(\tilde{a})) - p_i^X(\alpha(\tilde{a})) = \kappa_i(t)\big(\tau_i(t) - \pi^r(\tilde{a})\big).$$

Subtracting gives

$$t\Delta q_i - \Delta p_i = \kappa_i(t)\big(\pi^r(\tilde{a}) - \pi^r(a)\big),$$

where Δq_i and Δp_i are constants independent of t. Thus κ_i is affine in t. Since $\kappa_i(t) > 0$ for every $t \in \mathbb{R}$, it must be constant: $\kappa_i(t) \equiv c_i > 0$. Therefore, whenever bidder i strictly wins in kPA(r) at profile a,

$$tq_i^X(\alpha(a)) - p_i^X(\alpha(a)) = c_i\tau_i(t) - c_i\pi^r(a). \tag{7}$$

The left side is affine in t, so τ_i is affine. Since τ_i is a bijection of $\mathbb{R}$, we may write

$$\tau_i(t) = m_i t + n_i \qquad \text{with } m_i \neq 0.$$

Step 4: Winners in kPA(r) receive the object with probability one in X, giving $c_i m_i = 1$. If bidder i strictly wins in kPA(r) at a, then every other bidder has allocation probability zero in kPA(r). By (6), every other bidder also has allocation probability zero in X at $\alpha(a)$. Since $\alpha_i(\varnothing) = \varnothing$ and α_i is injective, bidder i's image action is not $\varnothing$. Hence the outside-option property implies that bidder i receives the object with probability one in X at $\alpha(a)$. Comparing coefficients in (7) gives

$$c_i m_i = 1, \qquad p_i^X(\alpha(a)) = c_i\big(\pi^r(a) - n_i\big) \tag{8}$$

whenever i strictly wins in kPA(r).

Step 5: The action relabelings coincide across bidders, i.e. $\alpha_i = \alpha_j$ for all i, j. Suppose, to the contrary, that for some bidders i, j and some $x \in \mathcal{A}_0$, $\alpha_i^{-1}(x) \neq \alpha_j^{-1}(x)$. Since $\alpha_i(\varnothing) = \alpha_j(\varnothing) = \varnothing$, this can only happen for $x \neq \varnothing$. Consider the X-profile where bidders i and j both choose x, and all other bidders choose $\varnothing$. By symmetry, bidders i and j must have the same allocation probability. But in the corresponding reserve-auction profile, bidders i and j make different non-null bids. Therefore exactly one of them strictly wins. By (6) and the above, the image profile in X gives allocation probability one to the strict winner and zero to the other bidder. This contradicts symmetry of X. Hence $\alpha_i = \alpha_j$ for all i, j. Write the common bijection as α_0.

Step 6: The constants c_i and n_i are common across bidders. Consider two bidders i and j, and two strict bid profiles that are permutations of each other, one in which i wins and one in which j wins, with the same payment-relevant price π. Since X is symmetric and α_0 is common, the winner's payment in X must be the same in the two profiles. By (8),

$$c_i(\pi - n_i) = c_j(\pi - n_j).$$

Since this equality holds for at least two distinct relevant values of π, we get $c_i = c_j$ and $n_i = n_j$, and so there exist constants $c > 0$ and $n \in \mathbb{R}$ such that $c_i = c$ and $n_i = n$ for every i. By (8), $m_i = 1/c$ for every i.

Step 7: Constructing the strategic equivalence with kPA$(\bar{r})$. We now show that X is strategically equivalent to a kth-price auction with an appropriate reserve. Take any reserve-auction profile a and any bidder i. If $q_i^{k\mathrm{PA}(r)}(a) = 0$, then (6) gives

$$q_i^X(\alpha(a)) = p_i^X(\alpha(a)) = 0.$$

If $q_i^{k\mathrm{PA}(r)}(a) > 0$, then

$$U_i^{k\mathrm{PA}(r)}[\tau_i(t), a] = q_i^{k\mathrm{PA}(r)}(a)\big(\tau_i(t) - \pi^r(a)\big).$$

Using $\kappa_i(t) = c$, $\tau_i(t) = t/c + n$, and $\lambda_i(t) = 0$, equation (2) gives

$$t q_i^X(\alpha(a)) - p_i^X(\alpha(a)) = c q_i^{k\mathrm{PA}(r)}(a)\left(\frac{t}{c} + n - \pi^r(a)\right).$$

Therefore,

$$q_i^X(\alpha(a)) = q_i^{k\text{PA}(r)}(a), \qquad p_i^X(\alpha(a)) = c q_i^{k\text{PA}(r)}(a)\big(\pi^r(a) - n\big). \tag{9}$$

Since payments in X are nonnegative, (9) evaluated at the lowest possible price r implies $c(r - n) \geq 0$. We can then define:

$$\bar{r} := c(r - n) \in \mathbb{R}_+.$$

Now define a bijection $\beta_i : \{\varnothing\} \cup [\bar{r}, \infty) \to \mathcal{A}_0$ such that

$$\beta_i(\varnothing) = \varnothing, \qquad \beta_i(b) = \alpha_0\left(\frac{b}{c} + n\right) \quad \text{for } b \geq \bar{r}.$$

This is well-defined since $b \geq \bar{r} = c(r - n)$ implies $\frac{b}{c} + n \geq r$. Take any profile b in $k\text{PA}(\bar{r})$, and let

$$a_i = \begin{cases} \varnothing, & b_i = \varnothing, \\ b_i/c + n, & b_i \neq \varnothing. \end{cases}$$

The transformation $b_i \mapsto b_i/c + n$ is strictly increasing, so it preserves participants, highest-bidder sets, allocation probabilities, and the ranking of bids. Moreover, if $\bar{\pi}(b)$ is the payment-relevant kth price in $k\text{PA}(\bar{r})$, and $\pi^r(a)$ is the payment-relevant kth price in $k\text{PA}(r)$, then

$$\bar{\pi}(b) = c(\pi^r(a) - n).$$

By (6) and (9), for every bidder i,

$$q_i^X(\beta(b)) = q_i^{k\text{PA}(\bar{r})}(b) \quad \text{and} \quad p_i^X(\beta(b)) = p_i^{k\text{PA}(\bar{r})}(b).$$

Hence, for every bidder i, every type t_i, and every profile b,

$$U_i^X[t_i, \beta(b)] = U_i^{k\text{PA}(\bar{r})}[t_i, b].$$

Therefore X is strategically equivalent to $k\text{PA}(\bar{r})$.

Step 8: Converse. The converse is immediate: if X is strategically equivalent to $k\text{PA}(\bar{r})$ for some $\bar{r} \in \mathbb{R}_+$, then it is strategically similar to a kth-price auction with reserve, taking τ_i to be the identity map for every bidder i and setting all $\kappa_i \equiv 1, \lambda_i \equiv 0$.

A.6 Proof of Proposition 4

Define, for each bidder i,

$$\alpha_i(\varnothing) = \varnothing, \quad \alpha_i(b) = \frac{c'}{c} b \text{ for } b \in \mathbb{R}_+, \quad \text{and} \quad \tau_i(t) = \frac{c}{c'} t, \quad \kappa_i(t) = \frac{c'}{c}, \quad \lambda_i(t) = 0.$$

Each α_i is a bijection from $\{\varnothing\} \cup \mathbb{R}_+$ to itself, and each τ_i is a bijection from $\mathbb{R}$ to itself. Fix a bidder i, a type $t_i \in \mathbb{R}$, and an action profile a. The map α preserves the set of participants and the ranking of bids. Hence it preserves the allocation probabilities induced by the symmetric tie-breaking rule. Moreover, if $a_i \neq \varnothing$, then the kth-highest participating bid is multiplied by c'/c:

$$\alpha(a)^{(k)} = \frac{c'}{c} a^{(k)}.$$

If $a_i = \varnothing$, then bidder i's payoff is zero in both mechanisms, so

$$U_i^{k\text{EC}(c')}[t_i, \alpha(a)] = 0 = \frac{c'}{c} U_i^{k\text{EC}(c)}[\tau_i(t_i), a].$$

Now suppose $a_i \neq \varnothing$. Let $q_i(a)$ be bidder i's allocation probability in the profile a. Since α preserves allocation probabilities, $q_i(\alpha(a)) = q_i(a)$. Therefore

$$U_i^{k\text{EC}(c')}[t_i, \alpha(a)] = q_i(a)\left(t_i - \frac{c'}{c} a^{(k)}\right) - c' = \frac{c'}{c}\left[q_i(a)\left(\frac{c\, t_i}{c'} - a^{(k)}\right) - c\right] = \frac{c'}{c} U_i^{k\text{EC}(c)}[\tau_i(t_i), a].$$

Thus, for every bidder i, type t_i, and action profile a,

$$U_i^{k\text{EC}(c')}[t_i, \alpha(a)] = \kappa_i(t_i)\, U_i^{k\text{EC}(c)}[\tau_i(t_i), a] + \lambda_i(t_i),$$

with $\kappa_i(t_i) = c'/c > 0$ and $\lambda_i(t_i) = 0$. Hence $k\text{EC}(c)$ and $k\text{EC}(c')$ are strategically similar.

A.7 Proof of Theorem 3

Proof. For mechanism $X^r = (\mathcal{A}^r, \Phi^r)$, write $q_i^r(a)$ for bidder i's probability of receiving the object at action profile a, and $p_i^r(a)$ her expected payment at this action profile.

Suppose toward a contradiction that $\mathcal{C}$ implements the Myerson-optimal auction for every symmetric IPV prior in symmetric equilibrium. Choose a symmetric IPV prior F with full support on $[\underline{\theta}, \overline{\theta}]$ whose ironed virtual value has a single ironing interval

$$I = [\theta_\ell, \theta_h], \qquad \underline{\theta} < \theta_\ell < \theta_h < \overline{\theta},$$

and whose ironed virtual value is strictly negative on a nonempty lower interval and strictly positive from some type in $(\underline{\theta}, \theta_\ell)$ onwards. Let $\sigma : \mathbb{R} \to \Delta(\mathcal{A}_i^1)$ be the symmetric equilibrium of X^1 implementing the Myerson-optimal auction for this prior.

Since X^0 and X^1 are strategically similar, there exist bijections

$$\alpha_i : \mathcal{A}_i^1 \to \mathcal{A}_i^0, \qquad \tau_i : \mathbb{R} \to \mathbb{R},$$

and functions

$$\kappa_i : \mathbb{R} \to \mathbb{R}_{++}, \qquad \lambda_i : \mathbb{R} \to \mathbb{R},$$

such that, for every bidder i, every type $\theta \in \mathbb{R}$, and every pure action profile $a \in \mathcal{A}^1$,

$$U_i^{X^0}[\tau_i(\theta), \alpha(a)] = \kappa_i(\theta) U_i^{X^1}[\theta, a] + \lambda_i(\theta). \tag{10}$$

Equivalently,

$$\tau_i(\theta) q_i^0(\alpha(a)) - p_i^0(\alpha(a)) = \kappa_i(\theta)\big(\theta q_i^1(a) - p_i^1(a)\big) + \lambda_i(\theta). \tag{11}$$

We first show three lemmas.

Lemma 1. *For every bidder i,*

$$q_i^0(\alpha(a)) = q_i^1(a) \quad \textit{for all } a \in \mathcal{A}^1.$$

Proof. Fix bidder i. We first show there exist types $\theta_0, \dots, \theta_3$ and constants $\gamma_0, \dots, \gamma_3$ such that

$$\sum_{r=0}^{3} \gamma_r = 0, \qquad \sum_{r=0}^{3} \gamma_r \kappa_i(\theta_r) = 0, \qquad \sum_{r=0}^{3} \gamma_r \kappa_i(\theta_r)\theta_r \neq 0.$$

If not, then every linear relation annihilating the functions $\theta \mapsto 1$ and $\theta \mapsto \kappa_i(\theta)$ would also annihilate the function $\theta \mapsto \kappa_i(\theta)\theta$. Hence, there would exist constants c_0^i, c_1^i such that

$$\kappa_i(\theta)\theta = c_0^i + c_1^i \kappa_i(\theta) \qquad \text{for all } \theta.$$

Setting $\theta = c_1^i$ gives $c_0^i = 0$. Therefore, for every $\theta \neq c_1^i$,

$$\kappa_i(\theta)(\theta - c_1^i) = 0,$$

which contradicts $\kappa_i(\theta) > 0$, so the desired types and constants exist.

We now apply (11) to each θ_r, multiply by γ_r, and sum over r. The payment terms vanish, and we obtain, for every action profile $a \in \mathcal{A}^1$,

$$\left(\sum_{r=0}^{3} \gamma_r \tau_i(\theta_r)\right) q_i^0(\alpha(a)) = \left(\sum_{r=0}^{3} \gamma_r \kappa_i(\theta_r)\theta_r\right) q_i^1(a) + \sum_{r=0}^{3} \gamma_r \lambda_i(\theta_r).$$

Let

$$S_i^0 := \sum_{r=0}^{3} \gamma_r \tau_i(\theta_r), \qquad S_i^1 := \sum_{r=0}^{3} \gamma_r \kappa_i(\theta_r)\theta_r, \qquad S_i^\lambda := \sum_{r=0}^{3} \gamma_r \lambda_i(\theta_r).$$

Then the preceding equality can be written as

$$S_i^0 q_i^0(\alpha(a)) = S_i^1 q_i^1(a) + S_i^\lambda \qquad \text{for all } a \in \mathcal{A}^1. \tag{12}$$

By construction, $S_i^1 \neq 0$. We also have $S_i^0 \neq 0$. This is because $S_i^0 = 0$ would imply

$$S_i^1 q_i^1(a) + S_i^\lambda = 0 \qquad \text{for all } a \in \mathcal{A}^1,$$

and, since X^1 implements the Myerson allocation for a nondegenerate full-support prior, there are equilibrium action profiles $\hat{a}^0, \hat{a}^1$ at which bidder i's allocation probabilities are respectively 0 and 1. Evaluating the preceding display at these two profiles gives $S_i^\lambda = 0$ and $S_i^1 + S_i^\lambda = 0$, so $S_i^1 = 0$; contradiction.

We first show that $|S_i^0| \leq |S_i^1|$. Since X^0 has a highest-bidder-wins allocation rule and each α_j is onto, there exist action profiles $a^0, a^1 \in \mathcal{A}^1$ such that $q_i^0(\alpha(a^0)) = 0$ and $q_i^0(\alpha(a^1)) = 1$. Evaluating (12) at a^1 and a^0, and then subtracting, gives

$$S_i^0 = S_i^1 \left(q_i^1(a^1) - q_i^1(a^0) \right).$$

Hence

$$|S_i^0| = |S_i^1| \left| q_i^1(a^1) - q_i^1(a^0) \right| \leq |S_i^1|,$$

because $q_i^1(a) \in [0,1]$.

Conversely, since X^1 implements the Myerson allocation for a nondegenerate full-support prior, there are equilibrium action profiles $\hat{a}^0, \hat{a}^1$ at which bidder i's allocation probabilities are respectively 0 and 1. Subtracting the corresponding equalities gives

$$|S_i^1| = |S_i^0| \left| q_i^0(\alpha(\hat{a}^1)) - q_i^0(\alpha(\hat{a}^0)) \right| \leq |S_i^0|,$$

because $q_i^0(\alpha(a)) \in [0,1]$. Therefore $|S_i^0| = |S_i^1|$. If $S_i^0 = S_i^1$, then

$$q_i^0(\alpha(a)) = q_i^1(a) + \frac{S_i^\lambda}{S_i^0}.$$

Since $q_i^0(\alpha(a))$ takes both values 0 and 1, while $q_i^1(a) \in [0,1]$, this implies $S_i^\lambda / S_i^0 = 0$. Hence

$$q_i^0(\alpha(a)) = q_i^1(a) \qquad \text{for all } a \in \mathcal{A}^1.$$

If $S_i^0 = -S_i^1$, then

$$q_i^0(\alpha(a)) = -q_i^1(a) + \frac{S_i^\lambda}{S_i^0}.$$

Again using that $q_i^0(\alpha(a))$ takes both values 0 and 1, while $q_i^1(a) \in [0,1]$, we get $S_i^\lambda / S_i^0 = 1$. Hence

$$q_i^0(\alpha(a)) = 1 - q_i^1(a) \qquad \text{for all } a \in \mathcal{A}^1. \tag{13}$$

We now rule out the alternative (13). Choose a type θ^0 whose ironed virtual value is strictly negative, and choose a type $\theta^+ > \theta_h$ whose ironed virtual value is strictly positive. Let

$$C := \left\{ i \in \mathcal{I} : q_i^0(\alpha(a)) = 1 - q_i^1(a) \text{ for all } a \in \mathcal{A}^1 \right\}.$$

We show that $C = \varnothing$.

Suppose first that $C = \mathcal{I}$. At the type profile $(\theta^0, \dots, \theta^0)$, every bidder has strictly negative ironed virtual value, so the Myerson-optimal auction does not allocate the object. Since X^1 implements this allocation and allocation probabilities are nonnegative,

$$q_j^1(a) = 0 \qquad \text{for every } j \in \mathcal{I}$$

for almost every action profile a induced by equilibrium play at this type profile. The complement alternative therefore gives $q_j^0(\alpha(a)) = 1$ for every $j \in \mathcal{I}$ for almost every such action profile a, contradicting feasibility of X^0, since the object cannot be allocated with probability one to more than one bidder.

Now suppose C is nonempty but $C \neq \mathcal{I}$. Choose $k \in C$ and $i \notin C$. Since $i \notin C$, bidder i satisfies the direct alternative:

$$q_i^0(\alpha(a)) = q_i^1(a) \qquad \text{for all } a \in \mathcal{A}^1.$$

Consider the type profile at which bidder i has type θ^+ and all other bidders have type θ^0. At this profile, bidder i is the unique bidder with strictly positive ironed virtual value. Thus the Myerson-optimal auction gives bidder i the object with probability one and gives bidder k probability zero. Since X^1 implements this allocation,

$$q_i^1(a) = 1 \qquad \text{and} \qquad q_k^1(a) = 0$$

for almost every action profile a induced by equilibrium play at this type profile. The direct alternative for i and the complement alternative for k imply

$$q_i^0(\alpha(a)) = 1 \qquad \text{and} \qquad q_k^0(\alpha(a)) = 1$$

for almost every such action profile a, again contradicting feasibility of X^0. Hence $C = \varnothing$, proving the lemma. □

Recall that $A_{-i}(\sigma, F^N)$ denotes the distribution over opponents' actions induced by σ_{-i} under the prior F^N (since types are drawn independently, we suppress the conditioning on θ_i).

Lemma 2. *Fix a bidder i. If $\theta, \theta' \in I$, $a_i \in \operatorname{supp} \sigma_i(\theta)$, and $a_i' \in \operatorname{supp} \sigma_i(\theta')$, then*

$$\mathbb{E}_{\tilde{a}_{-i} \sim A_{-i}(\sigma, F^N)} \left[q_i^1(a_i, \tilde{a}_{-i}) \right] = \mathbb{E}_{\tilde{a}_{-i} \sim A_{-i}(\sigma, F^N)} \left[q_i^1(a_i', \tilde{a}_{-i}) \right].$$

Consequently,

$$\mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[q_i^0(\alpha_i(a_i),\alpha_{-i}(\tilde{a}_{-i}))\right] = \mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[q_i^0(\alpha_i(a_i'),\alpha_{-i}(\tilde{a}_{-i}))\right].$$

Proof. For $a_i \in \mathcal{A}_i^1$, define

$$Q_i^1(a_i) := \mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[q_i^1(a_i,\tilde{a}_{-i})\right] \quad \text{and} \quad P_i^1(a_i) := \mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[p_i^1(a_i,\tilde{a}_{-i})\right].$$

If $a_i \in \operatorname{supp}\sigma_i(\theta)$ and $a_i' \in \operatorname{supp}\sigma_i(\theta')$, optimality gives

$$\theta Q_i^1(a_i) - P_i^1(a_i) \geq \theta Q_i^1(a_i') - P_i^1(a_i') \quad \text{and} \quad \theta' Q_i^1(a_i') - P_i^1(a_i') \geq \theta' Q_i^1(a_i) - P_i^1(a_i).$$

Adding,

$$(\theta' - \theta)\big(Q_i^1(a_i') - Q_i^1(a_i)\big) \geq 0. \tag{14}$$

Thus support actions used by higher types cannot yield lower interim allocation.

Now, revenue-optimality requires every type in the ironing interval I to get the same interim allocation. Since the average interim allocation induced by $\sigma_i(\theta)$ is constant over $\theta \in I$, the monotonicity (14) implies that every support action used by every type in I yields this same interim allocation.

The second claim follows from Lemma 1. For each bidder i, the relabeled allocation in X^0 is $q_i^0 \circ \alpha$, so the equality is preserved. □

Lemma 3. *There exists an action $a_I \in \mathcal{A}_i^1$ such that, for every bidder i,*

$$\sigma_i(\theta) = \delta_{a_I} \qquad \text{for } F\text{-almost every } \theta \in I.$$

Moreover, the relabeled bid is common across bidders:

$$\alpha_i(a_I) = \alpha_j(a_I) \qquad \text{for all } i,j \in \mathcal{I}.$$

Proof. Define

$$Q_i^0(b_i) := \mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[q_i^0(b_i,\alpha_{-i}(\tilde{a}_{-i}))\right] \quad \text{and} \quad P_i^0(b_i) := \mathbb{E}_{\tilde{a}_{-i}\sim A_{-i}(\sigma,F^N)}\left[p_i^0(b_i,\alpha_{-i}(\tilde{a}_{-i}))\right].$$

For each bidder i, let μ_i^I be the distribution of bidder i's relabeled bid conditional on her type lying in I:

$$\mu_i^I := \int_I \alpha_{i\#}\sigma_i(\theta)\, dF(\theta \mid I).$$

By Lemmas 1 and 2, there is a number $\bar{Q}_i \in (0,1)$ such that

$$Q_i^0(b_i) = \bar{Q}_i \qquad \text{for } \mu_i^I\text{-almost every } b_i.$$

We show that each μ_i^I is a point mass. Suppose not. Then, for some bidder i, we can choose

$$b_i^- < b_i^+$$

in the essential support of μ_i^I, at points where $Q_i^0(b_i^-) = Q_i^0(b_i^+) = \bar{Q}_i$.

Since X^0 gives the good to the highest participating bidder and breaks ties uniformly, raising bidder i's relabeled bid from b_i^- to b_i^+ strictly raises her allocation probability if the highest competing bid lies in $[b_i^-, b_i^+]$ with positive probability. But b_i^- and b_i^+ give bidder i the same interim allocation, so

$$\mathbb{P}_{\tilde{a}_{-i} \sim A_{-i}(\sigma, F^N)}\left(\max_{j \neq i: \alpha_j(\tilde{a}_j) \neq \varnothing} \alpha_j(\tilde{a}_j) \in [b_i^-, b_i^+] \right) = 0, \tag{15}$$

with the convention that the maximum equals $-\infty$ if no opponent participates.

This implies that no opponent's relabeled ironed-bid distribution can put positive mass on $[b_i^-, b_i^+]$. That is, for every $j \neq i$,

$$\mu_j^I\big([b_i^-, b_i^+]\big) = 0. \tag{16}$$

Indeed, since $Q_i^0(b_i^-) > 0$, bidder i wins with positive probability when bidding b_i^-. This implies that all opponents' relabeled bids are strictly below b_i^- with positive probability. If, for some $j \neq i$, bidder j's relabeled ironed bid lay in $[b_i^-, b_i^+]$ with positive probability, then, by independence of types and independent randomization, this would also occur with positive probability while all other opponents $l \neq i, j$ bid below b_i^-. On this event, the highest competing relabeled bid lies in $[b_i^-, b_i^+]$, contradicting (15).

Repeating the same argument with any bidder in the role of i gives the following no-crossing property: if $c_k^- < c_k^+$ are two essential support points of μ_k^I at which $Q_k^0 = \bar{Q}_k$, then

$$\mu_\ell^I([c_k^-, c_k^+]) = 0 \qquad \text{for every } \ell \neq k.$$

Thus the relabeled ironed-bid distributions of different bidders cannot interlace. Indeed, if some opponent $j \neq i$ put positive mass both below b_i^- and above b_i^+, then we could choose essential support points $c_j^- < b_i^- < b_i^+ < c_j^+$. Applying the no-crossing property to bidder j would give $\mu_i^I([c_j^-, c_j^+]) = 0$, contradicting that b_i^- and b_i^+ lie in the essential support of μ_i^I.

Consequently, for any bidder i with a nondegenerate μ_i^I, every other bidder's relabeled ironed-bid distribution lies either below μ_i^I almost surely or above μ_i^I almost surely.

Now define

$$E_I := \{\theta_j \in I \text{ for every } j \in \mathcal{I}\}.$$

This event has positive probability because I has positive F-measure. Conditional on E_I, the symmetric revenue-maximizing allocation gives each bidder probability $1/N$ of winning in mechanism X^1:

$$\mathbb{E}\left[q_i^1(a)\,\middle|\,E_I,\ a\sim\sigma(\theta)\right]=\tfrac{1}{N}\qquad\text{for every } i\in\mathcal{I}.$$

By Lemma 1, we then get

$$\mathbb{E}\left[q_i^0(\alpha(a))\,\middle|\,E_I,\ a\sim\sigma(\theta)\right]=\tfrac{1}{N}\qquad\text{for every } i\in\mathcal{I}. \tag{17}$$

Return to a bidder i with nondegenerate μ_i^I. By the ordering just established, either some opponent's relabeled ironed-bid distribution lies above μ_i^I almost surely, or every opponent's relabeled ironed-bid distribution lies below μ_i^I almost surely. In the first case, bidder i receives allocation 0 in the relabeled mechanism X^0 conditional on E_I. In the second case, bidder i receives allocation 1 conditional on E_I. Both conclusions contradict (17). Therefore μ_i^I must be a point mass for every bidder i.

Thus, for every bidder i, there is a relabeled bid b_i such that $\alpha_{i\#}\sigma_i(\theta)=\delta_{b_i}$ for F-almost every $\theta\in I$. Since each α_i is a bijection and the equilibrium strategy is symmetric in X^1, there exists an action $a_I\in\mathcal{A}_i^1$ such that $\sigma_i(\theta)=\delta_{a_I}$ for every bidder i and F-almost every $\theta\in I$.

It remains to show that the relabeled bid is common across bidders. Suppose not. Then there exists a bidder $\underline{i}$ such that $\alpha_{\underline{i}}(a_I)\le\alpha_j(a_I)$ for all $j\in\mathcal{I}$, with strict inequality for at least one j. Conditional on E_I, all bidders play a_I almost surely, so bidder $\underline{i}$'s relabeled bid is strictly below some opponent's relabeled bid. Hence bidder $\underline{i}$ receives allocation 0 in X^0 conditional on E_I, contradicting (17). Therefore $\alpha_i(a_I)=\alpha_j(a_I)$ for all $i,j\in\mathcal{I}$. □

We now finish the proof. Let

$$b_I:=\alpha_i(a_I),$$

which is independent of i by Lemma 3. We first show that $b_I>\underline{a}$. Note there is a type $\theta^-<\theta_\ell$ that wins with positive probability in the Myerson allocation, but loses whenever at least one opponent has type in I. Choose $a^-\in\operatorname{supp}\sigma_i(\theta^-)$ with positive interim allocation and set $b^-:=\alpha_i(a^-)$. Then $b^-\neq\varnothing$. If bidder i plays a^- while some opponent plays a_I, the Myerson allocation gives bidder i probability zero. By the allocation-orientation step, the corresponding relabeled profile in X^0 also gives bidder i probability zero. Since X^0 allocates to the highest participating bid and the opponent's relabeled bid is b_I, we must have $b^-<b_I$.

Now, fix a bidder i and a type $\theta\in I$. Let T_i be the event that bidder i's opponents' highest relabeled bid is exactly b_I:

$$T_i:=\left\{\max_{j\neq i:\,\alpha_j(\tilde a_j)\neq\varnothing}\alpha_j(\tilde a_j)=b_I\right\},$$

with the convention that the maximum is $-\infty$ if no opponent participates. Also define the tie-payment wedge

$$\Delta^0(b_I):=\pi_w^0(b_I,(b_I)_{j\neq i})-\pi_\ell^0(b_I,(b_I)_{j\neq i}).$$

Indeed, this difference only depends on b_I since X^0 satisfies payment regularity.

Consider first the upward deviation in X^1 to $\alpha_i^{-1}(b_I+\epsilon)$. Since type θ plays a_I in equilibrium, incentive compatibility in X^1, together with (10), implies

$$U_i^{X^0}\left[\tau_i(\theta),\delta_{b_I}\otimes\alpha_{-i\#}A_{-i}(\sigma,F^N\mid\theta)\right]\geq U_i^{X^0}\left[\tau_i(\theta),\delta_{b_I+\epsilon}\otimes\alpha_{-i\#}A_{-i}(\sigma,F^N\mid\theta)\right].\tag{18}$$

Letting $\epsilon\downarrow 0$, all profiles outside T_i have vanishing contribution. Indeed, if the highest opposing relabeled bid is strictly below b_I, bidder i wins both before and after the deviation, and continuity and local boundedness of π_w^0 imply that the payment difference vanishes. If the highest opposing relabeled bid is strictly above b_I, bidder i loses both before and after the deviation, except on the event that the highest opposing bid lies in $(b_I,b_I+\epsilon]$, whose probability tends to zero; local boundedness of π_ℓ^0 makes this contribution vanish.

On T_i, the upward deviation changes bidder i's allocation from $q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))$ to 1. Using the tie-payment assumption, the limiting contribution on T_i is therefore

$$\mathbb{1}_{T_i}\left(1-q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))\right)\left(\Delta^0(b_I)-\tau_i(\theta)\right).$$

Taking limits in (18) gives

$$0\leq\mathbb{E}_{a_{-i}\sim A_{-i}(\sigma,F^N)}\left[\mathbb{1}_{T_i}\left(1-q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))\right)\right]\left(\Delta^0(b_I)-\tau_i(\theta)\right).$$

The expectation multiplying the parentheses is strictly positive: with positive probability all opponents have types in I, and then, by Lemma 3, all opponents play a_I, so all their relabeled bids equal b_I. On that event bidder i's allocation from bidding b_I is $1/N<1$. Hence

$$\tau_i(\theta)\leq\Delta^0(b_I).$$

Now consider the downward deviation in X^1 to $\alpha_i^{-1}(b_I-\epsilon)$ for $\epsilon>0$ small enough that $b_I-\epsilon\in[\underline{a},\infty)$. The same argument gives

$$U_i^{X^0}\left[\tau_i(\theta),\delta_{b_I}\otimes\alpha_{-i\#}A_{-i}(\sigma,F^N\mid\theta)\right]\geq U_i^{X^0}\left[\tau_i(\theta),\delta_{b_I-\epsilon}\otimes\alpha_{-i\#}A_{-i}(\sigma,F^N\mid\theta)\right].\tag{19}$$

Again, after letting $\epsilon\downarrow 0$, only profiles in T_i have a nonvanishing contribution. On T_i, the downward deviation changes bidder i's allocation from $q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))$ to 0. Thus the limiting contribution is

$$\mathbb{1}_{T_i}q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))\left(\tau_i(\theta)-\Delta^0(b_I)\right).$$

Taking limits in (19) gives

$$0\leq\mathbb{E}_{a_{-i}\sim A_{-i}(\sigma,F^N)}\left[\mathbb{1}_{T_i}q_i^0(b_I,\alpha_{-i}(\tilde{a}_{-i}))\right]\left(\tau_i(\theta)-\Delta^0(b_I)\right).$$

The expectation multiplying the parentheses is strictly positive: with positive probability all opponents have types in I, and on that event bidder i's allocation from bidding b_I is $1/N>0$.

Therefore

$$\tau_i(\theta) \geq \Delta^0(b_I).$$

Combining the two inequalities gives $\tau_i(\theta) = \Delta^0(b_I)$ for every $\theta \in I$. But I is nondegenerate and $\tau_i : \mathbb{R} \to \mathbb{R}$ is a bijection, so τ_i cannot be constant on I; contradiction. □

A.8 Proof of Proposition 5

We first consider the linear-score family. Fix $\lambda_1, \lambda_2 \in (0,1)$. For each seller i, define

$$\alpha_i(b_i) := b_i + \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) q_i, \qquad \tau_i(c_i) := c_i - \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) q_i, \qquad \kappa_i(c_i) := 1, \qquad \lambda_i(c_i) := 0.$$

Each $\alpha_i : \mathbb{R} \to \mathbb{R}$ and each $\tau_i : \mathbb{R} \to \mathbb{R}$ is a bijection.

For every bid profile b,

$$S^L_{\lambda_2}(\alpha(b))_i = -\lambda_2 \left(b_i + \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) q_i\right) + (1-\lambda_2) q_i = \frac{\lambda_2}{\lambda_1}\left[-\lambda_1 b_i + (1-\lambda_1) q_i\right] = \frac{\lambda_2}{\lambda_1} S^L_{\lambda_1}(b)_i.$$

Thus all sellers' scores are multiplied by the same positive constant. Hence the set of highest-scoring sellers, and therefore the winner selected by the fixed tie-breaking rule, is unchanged:

$$w_{L(\lambda_2)}(\alpha(b)) = w_{L(\lambda_1)}(b).$$

Therefore, for every seller i, every cost c_i, and every bid profile b,

$$\begin{aligned} U_i^{L(\lambda_2)}[c_i, \alpha(b)] &= \mathbb{1}\{w_{L(\lambda_1)}(b) = i\} \left(b_i + \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) q_i - c_i\right) \\ &= \mathbb{1}\{w_{L(\lambda_1)}(b) = i\} \left(b_i - \left(c_i - \left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right) q_i\right)\right) = U_i^{L(\lambda_1)}[\tau_i(c_i), b]. \end{aligned}$$

This is the affine strategic-similarity condition with $\kappa_i(c_i) = 1$ and $\lambda_i(c_i) = 0$. Hence $\{L(\lambda) : \lambda \in (0,1)\}$ is a class of strategically similar mechanisms.

Now consider the ratio-score family. Suppose the fixed quality scores are not all equal. Choose sellers 1 and 2 such that

$$q_1 = \max_j q_j \qquad \text{and} \qquad q_1 > q_2.$$

Choose $\lambda_1, \lambda_2 \in (0,1)$ such that

$$\lambda_1 < \frac{(q_1 - q_2)}{1 + (q_1 - q_2)} < \lambda_2. \tag{20}$$

We first show that seller 2 can never win in $R(\lambda_1)$. For any bid profile $b \in \mathbb{R}^N_{++}$, we have

$\min_j b_j/b_2 \leq 1$, and so seller 2's score is at most $\lambda_1 + (1-\lambda_1)q_2$. Because $\lambda_1 < \frac{(q_1-q_2)}{1+(q_1-q_2)}$, we have

$$\lambda_1 + (1-\lambda_1)q_2 < (1-\lambda_1)q_1.$$

Moreover, note that seller 1's score is strictly larger than $(1-\lambda_1)q_1$. Thus seller 2's score is always strictly below seller 1's score, and so seller 2 never wins in $R(\lambda_1)$.

Next we show that seller 2 can win in $R(\lambda_2)$. By (20), we then have:

$$\lambda_2 + (1-\lambda_2)q_2 > (1-\lambda_2)q_1.$$

Choose $b_2 > 0$; we can then make all b_j for $j \neq 2$ sufficiently large that $\min_k b_k/b_2 = 1$ and

$$\frac{\min_k b_k}{b_j} = \frac{b_2}{b_j}$$

is arbitrarily close to 0 for every $j \neq 2$. Hence, seller 2's score is $\lambda_2 + (1-\lambda_2)q_2$, while every other seller's score is arbitrarily close to

$$(1-\lambda_2)q_j \leq (1-\lambda_2)q_1.$$

Thus, for b_j large enough for all $j \neq 2$, seller 2 is the unique winner in $R(\lambda_2)$.

We now prove that $R(\lambda_1)$ and $R(\lambda_2)$ cannot be strategically similar. Suppose, toward a contradiction, that they are strategically similar. Then there exist bijections $\alpha_i : \mathbb{R}_{++} \to \mathbb{R}_{++}$, bijections $\tau_i : \mathbb{R} \to \mathbb{R}$, and functions $\kappa_i : \mathbb{R} \to \mathbb{R}_{++}$ and $\lambda_i : \mathbb{R} \to \mathbb{R}$ such that, for every seller i, every cost c_i, and every bid profile b,

$$U_i^{R(\lambda_2)}[c_i, \alpha(b)] = \kappa_i(c_i) U_i^{R(\lambda_1)}[\tau_i(c_i), b] + \lambda_i(c_i).$$

Apply this to seller 2. Since seller 2 never wins in $R(\lambda_1)$,

$$U_2^{R(\lambda_1)}[\tau_2(c_2), b] = 0 \qquad \text{for every } c_2 \text{ and every } b.$$

Therefore

$$U_2^{R(\lambda_2)}[c_2, \alpha(b)] = \lambda_2(c_2) \qquad \text{for every } c_2 \text{ and every } b.$$

Since α is onto, this implies that, for every fixed c_2, seller 2's payoff in $R(\lambda_2)$ is constant across all bid profiles, which is not the case. In $R(\lambda_2)$, seller 2 loses at some bid profiles, so her payoff is 0. She also wins at some bid profiles. Taking $c_2 = 0$ and a winning bid profile with $b_2 > 0$, her payoff is $b_2 > 0$. Thus her payoff is not constant across bid profiles, a contradiction.

A.9 Proof of Proposition 6

Input-based pricing. Fix two efficacy levels $e, e' > 0$. Write

$$\eta := \frac{e'}{e}.$$

For each buyer i, choose the action and type bijections

$$\alpha_i^{e \to e'}(a_i) := a_i, \qquad \tau_i^{e' \to e}(v_i)(x) := v_i(\eta x).$$

The former is a bijection because the action space is $[0,1]$ in every efficacy state. The latter is well-defined because if $v_i \in \mathcal{V}$ then $\tau_i^{e' \to e}(v_i) \in \mathcal{V}$, and is a bijection because it is invertible via

$$(\tau_i^{e' \to e})^{-1}(w_i)(x) = w_i(x/\eta).$$

Take any action profile $a \in [0,1]^{\mathcal{I}}$. Since the action relabeling is the identity, the requested capacities are unchanged. Therefore the rationing factor $\min\{1, 1/\sum_{j \in \mathcal{I}} a_j\}$ is unchanged, and so the allocated capacity

$$z_i(a) = a_i \min\left\{1, 1/\sum_{j \in \mathcal{I}} a_j\right\},$$

is the same under efficacy e and e'. Hence buyer i's output scales with efficacy:

$$x_i(\alpha^{e \to e'}(a), e') = e' z_i(a) = \eta\, e\, z_i(a) = \eta\, x_i(a, e).$$

Payments are unchanged because the tariff depends only on the allocated capacity:

$$p_i(\alpha^{e \to e'}(a)) = P(z_i(a)) = p_i(a).$$

Now fix buyer i, type $v_i \in \mathcal{V}$, and action profile $a \in [0,1]^{\mathcal{I}}$. Using the identities above,

$$U_i^{X_{e'}^{\text{in}}}[v_i, \alpha^{e \to e'}(a)] = v_i\left(x_i(\alpha^{e \to e'}(a), e')\right) - p_i(\alpha^{e \to e'}(a)) = v_i(\eta\, x_i(a, e)) - P(z_i(a)),$$

$$U_i^{X_e^{\text{in}}}[\tau_i^{e' \to e}(v_i), a] = \tau_i^{e' \to e}(v_i)\left(x_i(a, e)\right) - p_i(a) = v_i(\eta\, x_i(a, e)) - P(z_i(a)).$$

Therefore,

$$U_i^{X_{e'}^{\text{in}}}[v_i, \alpha^{e \to e'}(a)] = U_i^{X_e^{\text{in}}}[\tau_i^{e' \to e}(v_i), a].$$

Thus X_e^{in} and $X_{e'}^{\text{in}}$ are strategically similar with $\kappa_i(v_i) = 1$ and $\lambda_i(v_i) = 0$.

Output-based pricing. We show that $\{X_e^{\text{out}} : e > 0\}$ can be a class of strategically similar mechanisms only if $P(x) = Ax^{\rho}$ for some $A > 0$ and $\rho > 1$. Suppose it is a family of strategically similar mechanisms. Fix two efficacy levels $e, e' > 0$, and let X_e^{out} and $X_{e'}^{\text{out}}$ denote their corresponding mechanisms. Then there exist bijections $\alpha_i : [0, e] \to [0, e']$, $\tau_i : \mathcal{V} \to \mathcal{V}$ and functions $\kappa_i : \mathcal{V} \to \mathbb{R}_{++}$,

$\lambda_i : \mathcal{V} \to \mathbb{R}$ such that, for every buyer i, every $v_i \in \mathcal{V}$, and every $a \in [0,e]^N$,

$$U_i^{X_{e'}^{\text{out}}}[v_i, \alpha(a)] = \kappa_i(v_i)\, U_i^{X_e^{\text{out}}}[\tau_i(v_i), a] + \lambda_i(v_i). \tag{21}$$

For an action profile $a \in [0,e]^N$, write $x_i(a,e)$ for buyer i's allocation in efficacy state e.

We first show that $x_i(\alpha(a), e')$ depends on a only through $x_i(a,e)$, i.e. that there exists a function $h_i : [0,e] \to [0,e']$ such that

$$x_i(\alpha(a), e') = h_i\big(x_i(a,e)\big) \qquad \text{for every } a \in [0,e]^N. \tag{22}$$

Let $0 \in \mathcal{V}$ denote a constant zero valuation. Applying (21) to $v_i = 0$ gives

$$-P(x_i(\alpha(a), e')) = \kappa_i(0)\left(\tau_i(0)(x_i(a,e)) - P(x_i(a,e))\right) + \lambda_i(0).$$

The right-hand side depends on a only through $x_i(a,e)$. Since P is strictly increasing, it follows that $x_i(\alpha(a), e')$ also depends on a only through $x_i(a,e)$, establishing (22).

Fact 1. *$\alpha_i(0) = 0$ for every i.*

Proof. Fix i. If $a_i = 0$, then $x_i(a,e) = 0$ for every a_{-i}. Therefore, by (22), $x_i((\alpha_i(0), \alpha_{-i}(a_{-i})), e')$ is independent of a_{-i}. Suppose $\alpha_i(0) > 0$. Choose some $j \neq i$. For each $\ell \neq i, j$, choose $a_\ell \in [0,e]$ such that $\alpha_\ell(a_\ell) = 0$, which is possible because α_ℓ is onto. Now vary a_j so that $\alpha_j(a_j)$ ranges over all of $[0,e']$. Then buyer i's allocation under efficacy e' is

$$\alpha_i(0) \min\left\{1, \frac{e'}{\alpha_i(0) + \alpha_j(a_j)}\right\}.$$

This expression is not constant as $\alpha_j(a_j)$ varies from 0 to e', because $\alpha_i(0) > 0$. This contradicts the independence from a_{-i}. □

Now take a profile in which buyer i requests $x \in [0,e]$ and every other buyer requests 0. Since $\alpha_j(0) = 0$ for every $j \neq i$, buyer i's allocation under efficacy e' is $\alpha_i(x)$. Therefore $h_i(x) = \alpha_i(x)$. Thus (22) becomes

$$x_i(\alpha(a), e') = \alpha_i(x_i(a,e)) \qquad \text{for every } a \in [0,e]^N. \tag{23}$$

Normalize the action relabelings by defining

$$f_i(z) := \frac{\alpha_i(ez)}{e'}, \qquad z \in [0,1].$$

Each f_i is a bijection from $[0,1]$ to $[0,1]$, and $f_i(0) = 0$.

Fact 2. *Each f_i is strictly increasing.*

Proof. Fix $0 \le a < b \le 1$. If $a = 0$, then $f_i(a) = 0 < f_i(b)$, since f_i is injective and $f_i(0) = 0$. So suppose $a > 0$. Choose $j \ne i$ and choose

$$c \in (1 - b, 1 - a).$$

Consider profiles in which only buyers i and j make positive requests.

At normalized requests (a, c), total demand is below capacity, so buyer i's output under efficacy e is ea. By (23), the relabeled allocation under efficacy e' must be $\alpha_i(ea)$. Since $a > 0$, we have $\alpha_i(ea) > 0$. If the relabeled profile were rationed, buyer i's allocation would be strictly below his relabeled request $\alpha_i(ea)$, a contradiction. Hence the relabeled profile is not rationed, and therefore $f_i(a) + f_j(c) \le 1$.

At normalized requests (b, c), total demand exceeds capacity, so buyer i's output under e is

$$e\frac{b}{b+c} < eb.$$

By (23), the relabeled allocation under efficacy e' must be

$$\alpha_i\left(e\frac{b}{b+c}\right).$$

If the relabeled profile were not rationed, buyer i's allocation would be his relabeled request $\alpha_i(eb)$. But $e\frac{b}{b+c} < eb$ and α_i is injective, so $\alpha_i\left(e\frac{b}{b+c}\right) \ne \alpha_i(eb)$. Hence the relabeled profile must be rationed, and therefore $f_i(b) + f_j(c) > 1$. Combining the two inequalities gives $f_i(a) < f_i(b)$. □

Since f_i is a strictly increasing bijection from $[0,1]$ onto $[0,1]$, it is continuous and satisfies $f_i(1) = 1$. Now, fix distinct buyers $i \ne j$ and note two implications of (23). First, if $a, b \in [0,1]$ and $a + b = 1$, then

$$f_i(a) + f_j(b) = 1. \tag{24}$$

Indeed, if $a, b \in (0,1)$, the source profile is unrationed, so the relabeled profile is unrationed and $f_i(a) + f_j(b) \le 1$. For every small $\varepsilon > 0$, the source profile $(a + \varepsilon, b)$ is rationed, so the relabeled profile is rationed and $f_i(a + \varepsilon) + f_j(b) > 1$. Letting $\varepsilon \downarrow 0$ gives $f_i(a) + f_j(b) \ge 1$. Thus equality holds. The endpoint cases $a = 0$ or $b = 0$ follow from $f_i(0) = 0$ and $f_i(1) = 1$.

Second, if $a, b > 0$ and $a + b > 1$, then proportional rationing and (23) imply

$$f_i\left(\frac{a}{a+b}\right) = \frac{f_i(a)}{f_i(a) + f_j(b)}. \tag{25}$$

Fact 3. *There is a common increasing bijection $f : [0,1] \to [0,1]$ such that $f_i = f$ for every i, with $f(z) + f(1-z) = 1$ for every $z \in [0,1]$.*

Proof. By (24), we have $f_i(z) + f_j(1-z) = 1$ for every $z \in [0,1]$. Taking $a = b = s/2$ in (25), where

$s \in (1,2]$, gives

$$f_i(1/2) = \frac{f_i(s/2)}{f_i(s/2) + f_j(s/2)} \quad \Longrightarrow \quad \frac{f_i(s/2)}{f_j(s/2)} = \frac{f_i(1/2)}{1 - f_i(1/2)}.$$

At $s = 2$, the left-hand side equals 1, because $f_i(1) = f_j(1) = 1$. Therefore

$$f_i(z) = f_j(z) \qquad \text{for every } z \in [1/2, 1].$$

Using $f_i(z) + f_j(1-z) = 1$ and $f_j(z) + f_i(1-z) = 1$, this equality extends to all $z \in [0,1]$. □

Fact 4. *$f(z) = z$ for every $z \in [0,1]$.*

Proof. With a common f, equation (25) implies that, for all $a, b > 0$ with $a + b > 1$,

$$\frac{f\left(\frac{a}{a+b}\right)}{1 - f\left(\frac{a}{a+b}\right)} = \frac{f(a)}{f(b)}.$$

Now take $a = zy$ and $b = y$, where $z, y \in (0,1]$ and $y(1+z) > 1$. Then

$$\frac{f\left(\frac{z}{1+z}\right)}{1 - f\left(\frac{z}{1+z}\right)} = \frac{f(zy)}{f(y)} \quad \Longrightarrow \quad \frac{f\left(\frac{z}{1+z}\right)}{1 - f\left(\frac{z}{1+z}\right)} = f(z).$$

Therefore

$$f(zy) = f(z)f(y) \tag{26}$$

whenever $z, y \in (0,1]$ and $y(1+z) > 1$. This extends to all $z, y \in (0,1]$. Fix $z, y \in (0,1]$. Choose m large enough that

$$y^{1/m}(1 + zy) > 1.$$

Then, for each $\ell = 0, \ldots, m-1$,

$$y^{1/m}\left(1 + zy^{\ell/m}\right) > 1,$$

so (26) gives

$$f\left(zy^{(\ell+1)/m}\right) = f\left(zy^{\ell/m}\right) f\left(y^{1/m}\right).$$

Iterating,

$$f(zy) = f(z)f(y^{1/m})^m.$$

Applying the same argument with $z = 1$ gives $f(y) = f(y^{1/m})^m$. Hence

$$f(zy) = f(z)f(y) \qquad \text{for all } z, y \in (0,1].$$

Thus, the function $g(t) := \log f(e^t)$ for $t \le 0$ is continuous and additive:

$$g(s+t) = g(s) + g(t) \qquad \text{for all } s, t \le 0.$$

Therefore there exists $\rho > 0$ such that $f(z) = z^\rho$ for every $z \in (0,1]$. Since $f(z) + f(1-z) = 1$ for every $z \in [0,1]$, evaluating at $z = 1/2$ gives $2(1/2)^\rho = 1$, and so $\rho = 1$. □

Undoing the normalization then gives $\alpha_i(a_i) = \eta a_i$ for every i and $a_i \in [0,e]$, where $\eta := \frac{e'}{e}$, so

$$x_i(\alpha(a), e') = \eta\, x_i(a,e) \qquad \text{for every } i \text{ and every } a \in [0,e]^N. \tag{27}$$

We now use the convexity of P. Fix buyer i. Applying (21) to $v \equiv 0$ and using (27) gives

$$-P(\eta x) = \kappa_i(0)\big(\tau_i(0)(x) - P(x)\big) + \lambda_i(0) \qquad \text{for every } x \in [0,e].$$

At $x = 0$, we get $\lambda_i(0) = 0$. Therefore, for $c := \frac{1}{\kappa_i(0)} > 0$, we have

$$\tau_i(0)(x) = P(x) - cP(\eta x) \qquad \text{for every } x \in [0,e]. \tag{28}$$

Since $\tau_i(0) \in \mathcal{V}$, the function $x \mapsto P(x) - cP(\eta x)$ is weakly increasing and weakly concave on $[0,e]$. Because τ_i is onto, there exists $\bar{v}_i \in \mathcal{V}$ such that $\tau_i(\bar{v}_i) = 0$. Applying (21) to $\bar{v}_i$, again using (27), gives

$$\bar{v}_i(\eta x) - P(\eta x) = -\kappa_i(\bar{v}_i)P(x) + \lambda_i(\bar{v}_i) \qquad \text{for every } x \in [0,e].$$

At $x = 0$, we get $\lambda_i(\bar{v}_i) = 0$. Therefore, for $d := \kappa_i(\bar{v}_i) > 0$, we have

$$\bar{v}_i(\eta x) = P(\eta x) - dP(x) \qquad \text{for every } x \in [0,e]. \tag{29}$$

Since $\bar{v}_i \in \mathcal{V}$, the function $x \mapsto P(\eta x) - dP(x)$ is weakly increasing and weakly concave on $[0,e]$.

We use the following fact.

Fact 5. *Let $P : \mathbb{R}_+ \to \mathbb{R}_+$ be a regular price scheme. Fix $\eta > 0$, $K > 0$. If there exist $c, d > 0$ such that both*

$$x \mapsto P(x) - cP(\eta x) \qquad \textit{and} \qquad x \mapsto P(\eta x) - dP(x)$$

are weakly increasing and weakly concave on $[0,K]$, then

$$P(\eta x) = dP(x) \qquad \textit{for every } x \in [0,K].$$

Proof. Let $p(x)$ denote the right derivative of P at $x > 0$. Since P is convex, p exists and is weakly increasing. Since P is strictly increasing and convex, $p(x) > 0$ for every $x > 0$. Define

$$F(x) := P(x) - cP(\eta x), \qquad G(x) := P(\eta x) - dP(x).$$

By assumption, F and G are weakly increasing and weakly concave on $[0,K]$. Hence their right derivatives are nonnegative on $(0,K)$. Thus

$$F'_+(x) = p(x) - c\eta p(\eta x) \geq 0, \qquad G'_+(x) = \eta p(\eta x) - dp(x) \geq 0.$$

Combining these inequalities gives $p(x) \geq cd\, p(x)$. Since $p(x) > 0$, we get $cd \leq 1$.

Now use concavity. Since F and G are concave, their right derivatives are weakly decreasing. Hence, for any $0 < a < b < K$,

$$F'_+(b) \leq F'_+(a) \quad \Longrightarrow \quad p(b) - p(a) \leq c\eta\big(p(\eta b) - p(\eta a)\big).$$

Similarly,

$$G'_+(b) \leq G'_+(a), \quad \Longrightarrow \quad \eta\big(p(\eta b) - p(\eta a)\big) \leq d\big(p(b) - p(a)\big).$$

Combining the last two inequalities yields

$$p(b) - p(a) \leq cd\big(p(b) - p(a)\big).$$

Because P is strictly convex, p is not constant on $(0, K)$. Hence there exist $0 < a < b < K$ such that $p(b) > p(a)$. Therefore $cd \geq 1$. Together with $cd \leq 1$, this gives $cd = 1$.

Returning to the derivative inequalities, we have $p(x) \geq c\eta p(\eta x)$ and, since $d = 1/c$,

$$\eta p(\eta x) \geq \frac{1}{c} p(x).$$

Therefore $p(x) = c\eta p(\eta x)$ for every $x \in (0, K)$. Thus $F'_+(x) = 0$ for every $x \in (0, K)$. Since F is concave, this implies that F is constant on $[0, K]$. Because $F(0) = P(0) - cP(0) = 0$, we get

$$P(x) = cP(\eta x) \qquad \text{for every } x \in [0, K].$$

Since $cd = 1$, this is equivalent to

$$P(\eta x) = dP(x) \qquad \text{for every } x \in [0, K].$$

□

Applying this fact with $K = e$, we obtain a constant $C_{\eta,e} > 0$ such that

$$P(\eta x) = C_{\eta,e}\, P(x) \qquad \text{for every } x \in [0, e].$$

Because $e > 0$ was arbitrary, for every $\eta > 0$ there exists a constant $C(\eta) > 0$ such that

$$P(\eta x) = C(\eta)\, P(x) \qquad \text{for every } x \geq 0. \tag{30}$$

Indeed, the constants obtained from different values of e agree on overlapping intervals because $P(x) > 0$ for every $x > 0$. Setting $x = 1$ in (30) gives $C(\eta) = \frac{P(\eta)}{P(1)}$, and therefore, for every $x, \eta > 0$,

$$P(\eta x) = \frac{P(\eta)P(x)}{P(1)}.$$

Define $Q(x) := \frac{P(x)}{P(1)}$. Then $Q : \mathbb{R}_{++} \to \mathbb{R}_{++}$ is continuous, strictly increasing, and satisfies

$$Q(\eta x) = Q(\eta)Q(x) \qquad \text{for every } x, \eta > 0.$$

Thus Q is a continuous multiplicative function on $\mathbb{R}_{++}$. Hence there exists $\rho \in \mathbb{R}$ such that $Q(x) = x^{\rho}$. Since Q is strictly increasing, $\rho > 0$. Therefore $P(x) = P(1)\, x^{\rho}$. Writing $A := P(1) > 0$, we get

$$P(x) = Ax^{\rho}$$

for some $A > 0$ and $\rho > 0$. That is, P must be isoelastic. Strict convexity further requires $\rho > 1$. Therefore, for a generic regular price scheme, the output-pricing family $\{X_e^{\text{out}} : e > 0\}$ does not generate a class of strategically similar mechanisms.