Buying from a Group

Nima Haghpanah[*] Aditya Kuvalekar[†] Elliot Lipnowski[‡§]
Penn State *Essex* *Columbia*

March 27, 2024

Abstract

A buyer procures a good owned by a group of sellers whose heterogeneous cost of trade is private information. The buyer must either buy the whole good or nothing, and sellers share the transfer in proportion to their share of the good. We characterize the optimal mechanism: trade occurs if and only if the buyer's benefit of trade exceeds a weighted average of sellers' virtual costs. These weights are endogenous, with sellers who are ex ante less inclined to trade receiving higher weight. This mechanism always outperforms posted-price mechanisms. An extension characterizes the entire Pareto frontier.

JEL Classification: D23, D47, D71, D82, Q15

[*]haghpanah@psu.edu
[†]a.kuvalekar@essex.ac.uk
[‡]e.lipnowski@columbia.edu
[§]A previous version of this paper was circulated under the title "Selling to a Group." We thank Nageeb Ali, Dirk Bergemann, Dhruva Bhaskar, Ben Brooks, Yeon-Koo Che, Rahul Deb, Laura Doval, Alex Frankel, Marina Halac, Navin Kartik, George Mailath, Benny Moldovanu, Doron Ravid, Luis Rayo, Ilya Segal, and Ludvig Sinander for helpful discussions.

1. Introduction

For developing countries, a key challenge in transitioning from an agricultural economy to a manufacturing economy is land acquisition. Manufacturers often require large parcels of land whose ownership is dispersed among a group of individuals. Acquiring such land from a group of sellers is a challenging problem in the presence of property rights: no individual can be forced to part with his land.[1] In describing the puzzle of empty storefronts in prime areas in Moscow in the post-Soviet era, Heller (1998) terms such a situation "the tragedy of the anticommons": strong property rights lead to underuse of a resource. Several projects in India, for example, have sparked protests around issues of land acquisition largely because of unfair terms offered to the sellers or because they are coerced into selling against their wishes, resulting in those projects' relocation or complete abandonment.[2]

In a similar vein, one can think of redevelopment projects in cities. Redevelopment of apartment complexes typically involves a construction company building a larger apartment complex in place of an old one. The developer usually compensates the existing residents through apartments in the newly built building, where the apartment size is commensurate with the apartment size in the old building. Until recently, some Indian states, such as Gujarat, required the consent of *all* the residents of an apartment complex.

Or consider buying an indivisible asset, such as a business, from a group of siblings who have inherited it. The business has just one price, and owners typically receive a proportional share of the price. And as in the previous two examples, no person can be forced to agree to the offered terms of trade.[3]

While each of the above settings has its own idiosyncratic features, three common features stand out. First, a buyer wishes to purchase an indivisible good collectively owned by a group of sellers (agents). Each agent owns a fraction of the good, and while an agent's share is public information, his valuation of that share is private information. Second, strong property rights give any group mem-

[1] Sood (2020) argues that frictions associated with land fragmentation have hindered manufacturing growth in India. The effect of land fragmentation has also been studied in agriculture (Chand et al. (2011); Manjunatha et al. (2013)) and urban development (Gandhi et al. (2021)).

[2] One famous such case is that of protests in Singur in the state of West Bengal, India. The then government of West Bengal used eminent domain provisions to acquire 997 acres of land from farmers to allow Tata Motors to build a factory. The use of eminent domain for an arguably nonpublic project was met with massive protests that, eventually, led to the factory's being shifted out of West Bengal. See `https://en.wikipedia.org/wiki/Tata_Nano_Singur_controversy`.

[3] Kuran (2004) argues that Islamic inheritance laws have hindered the growth of Middle Eastern countries because they lead to fragmentation of enterprises and therefore prevent the creation of large-scale firms. Kuran documents that, by the nineteenth century, Western enterprises grew in size, but Middle Eastern enterprises did not. He suggests that Islamic inheritance laws played a role.

ber a refusal right—the right not to participate in trading. And third, the buyer can only offer one price for the entire good, and each seller receives a fraction of that price proportional to their ownership share.[4]

Motivated by such settings, we study the problem of acquiring a commonly owned good in a mechanism design setting with private information, voluntary participation, and ex-post transfers that are proportional to agents' shares. We assume that the buyer's valuation of the good is public information, but each seller's type is drawn independently from a publicly known distribution. In a (direct) mechanism, agents first report their types, and then the mechanism specifies the probability of trade and the price of trade (to be divided in proportion to agents' shares) as a function of the entire profile of reported types. Our main goal is to understand the buyer's profit-maximizing mechanism among all Bayesian incentive-compatible, interim individually rational mechanisms—henceforth the optimal mechanism.[5]

The first step toward understanding optimal mechanisms is to understand the class of implementable allocation rules in our setting. If the buyer could use transfers that need not be proportional to shares, standard arguments à la Myerson (1981) teach us that a given allocation rule would be implementable if and only if its associated interim allocations are nonincreasing. But in our setting, interim transfers are constrained because ex-post transfers must be proportional to shares, and it is a priori unclear what additional constraints this restriction places on the type of implementable mechanisms the buyer can offer. Even with this additional constraint, Lemma 1 shows that the same condition characterizes implementability in our setting. However, because of the proportional-transfers constraint, the agents' average per-share payments must coincide. Hence, the minimal average purchase price that can be attained for a given implementable allocation rule is pinned down by the condition that *one* agent's individual-rationality constraint is binding (and the others' are satisfied). Consequently, we can recast the buyer's problem as a maximin problem in which the maximum is over interim-monotone allocation rules and the minimum is over agents whose individual-rationality constraint is binding.

We solve the buyer's reformulated problem via an analogy to zero-sum games. We view the problem as a two-player zero-sum game in which one player (the

[4]For example, a recent land-acquisition bill in India stipulates that the landowners whose land has been acquired for private projects should be paid a specified fraction of the deemed market value of the land parcel. See `https://legislative.gov.in/sites/default/files/A2013-30.pdf` for details.

[5]The individual-rationality requirement, imposed for each agent, is meant to capture the strong property rights in the above examples. Our model imposes interim incentive constraints—that any agent must find it worthwhile in expectation (being uncertain of others' valuations) to participate in the mechanism and to report his valuation truthfully. In Section 7 we discuss ex post versions of both constraints.

Maximizer) chooses an allocation rule and the other player (the Minimizer) chooses an agent but may use a mixed strategy. We characterize the equilibria of this game to establish that the optimal allocation rule is unique and is a weighted allocation rule: the good is sold if and only if the buyer's benefit is larger than the weighted sum of agents' virtual valuations. These weights are endogenously determined, and are characterized by a simple program. These results are summarized in Theorem 1.

Given that the optimal mechanism assigns a weight to each agent, we study which agents are assigned higher weights. At a high level, agents with higher weights have more influence over the outcomes of the mechanism since trade is more sensitive to their reports. Theorem 2 answers this question by giving a condition under which we can rank agents' weights. In short, the optimal mechanism assigns higher weight to agents that have a higher valuation of the good in per-unit terms. More precisely, agent i has a higher valuation than agent j if i's virtual cost is higher in the *reversed hazard-rate order* than j's virtual cost.

The question of which agents have higher weights seems especially relevant to our land-acquisition application. Might the optimal mechanism discriminate against certain agents based on their characteristics? Our ranking result says that in this application, the optimal mechanism assigns more weight to agents with more productive land as measured in per-unit terms. Which agents are more productive depends on the context. Suppose first that there are two sellers who differ in how they use their land. Say agent 1 has a larger plot of land on which he can install a factory, while agent 2 has a smaller plot of land that he can use only for farming. Then, it is conceivable that (per square foot of land) agent 1 typically has higher productivity than agent 2. In this case, the optimal mechanism puts more weight on agent 1. In contrast, suppose that both the agents are small-scale farmers who differ only in their plot sizes. A literature in development economics has documented an inverse relationship between plot size and productivity (e.g., Banerjee, 1999). If this relationship were to hold for our two agents, then agent 1—the agent with a larger parcel of land—would have lower productivity (per square foot of land) than agent 2 and so would be granted a smaller weight.

Even though any mechanism in our setting offers all agents a uniform price per share, this price might depend on the entire profile of agents' types. But one might wonder whether the *optimal* mechanism could nonetheless be a posted-price mechanism in which the price is fixed. Such mechanisms are known to be optimal when there is one seller. However, with two or more sellers in the group, posted-price mechanisms are strictly suboptimal (Proposition 2). This result is derived from Theorem 1, which says that the price—even conditional on trade—must be responsive to the profile of reported types by the group members, a property evidently not shared by posted-price mechanisms.

In Section 6, we apply our analysis to shed light on the differences between mechanism design for a group versus an individual. To do so, we compare optimal mechanisms in our setting with optimal mechanisms in a benchmark setting in which a *single agent* owns the entire good. We first observe that optimal allocations in both settings have the familiar downward distortion to reduce information rent to the high types. However, trade outcomes are additionally distorted in the group setting in two novel ways. One distortion rotates the trade region in the type space, and the other affects its curvature. The overall effect depends on how the three types of distortions interact, and these distortions might even lead to a form of *over-trading* in which trade occurs when doing so is inefficient. We then compare the efficiency of optimal mechanisms in these two settings. We show that the ranking depends crucially on how large a benefit the good generates to the buyer. If this benefit is low, then the single-agent setting generates greater surplus; and if this benefit is high, then our group setting generates greater surplus.

While most of our analysis focuses on buyer-optimal mechanisms, one might consider alternative bargaining arrangements that give more power to the sellers. A more general notion of efficiency can be especially compelling when the buyer is a government or similar entity that might have a significant concern for nonmonetary welfare outcomes. With this in mind, in Theorem 3 we fully characterize the set of all the Pareto-optimal mechanisms. This characterization is facilitated by techniques similar to those we develop en route to Theorem 1. As we show, any Pareto-optimal mechanism allocates the good if and only if the weighted sum of agents' *actual and virtual costs* is lower than the benefit to the buyer.

We impose the assumption that the buyer must pay agents the same price per share as a fairness or institutional requirement that is natural in our main applications. We show additionally that this restriction reduces the agents' incentives to collude in a certain sense. In particular, we study a larger game in which the agents may trade their shares *before* interacting with the mechanism. We assume the sellers have identical distributions of value (i.e., their cost of trade) per unit but may possibly be endowed with nonidentical shares. We show that the sellers optimally choose not to trade shares if they are paid the same price per share, but they may benefit from trading if they are paid different prices per share. The literature on uniform- versus discriminatory-price auctions also argues that uniform-price auctions might reduce agents' incentives to collude (Friedman, 1960). Further, this literature points out that agents might be more likely to participate in uniform-price auctions (Malvey and Archibald, 1998; Ausubel et al., 2014).

1.1. Related work

Because the buyer procures the good from all sellers or none, our work is closely related to the literature on designing mechanisms for the provision of public goods. The canonical model (e.g., d'Aspremont and Gérard-Varet, 1979) allows for arbitrary monetary transfers between agents. Rob (1989) shows that with a large number of agents, profit-maximizing mechanisms are very inefficient, and Mailath and Postlewaite (1990) extend this inefficiency result to all incentive-feasible mechanisms. In a setting in which agents' values for a good are symmetric, and each is initially endowed with a share, Cramton et al. (1987) show efficient and individually rational trading mechanisms exist if and only if agents' shares are sufficiently symmetric. Ekmekci et al. (2016) identify profit-maximizing mechanisms for selling some fraction of a firm owned by a single agent to a single buyer. Güth and Hellwig (1986) identify profit-maximizing mechanisms for public good provision subject to incentive-compatibility and individual-rationality constraints. Hence, our buyer's problem is equivalent (up to a sign change) to that of Güth and Hellwig (1986), with the added restriction that transfers must be proportional to shares. Virtual costs (or values) often appear in the literature that studies profit-maximizing mechanisms, but a special feature of profit-maximizing mechanisms in our setting is that virtual costs are multiplied by endogenous weights that arise because of the proportional-transfers constraint.[6] These weights are interpretable as the degree of influence that the optimal mechanism gives to different agent, and we study how this influence is affected by seller heterogeneity.

Another strand of the literature on public goods studies voting mechanisms without monetary transfers. Starting with Rae (1969), many papers in this literature study mechanisms that maximize utilitarian efficiency. Schmitz and Tröger (2012) and Krishna and Morgan (2015) identify conditions under which a (weighted) majority does or does not maximize efficiency. Azrieli and Kim (2014) show any incentive-compatible mechanism must be a weighted-majority rule, and they characterize the weights that maximize efficiency.[7] The (weighted) majority structure of mechanisms in this literature is typically either assumed or is a property of all incentive compatible mechanisms. In our setting, on the other hand, the weights arise only as a feature of *optimal* mechanisms and are not necessarily a feature of all incentive compatible mechanisms.

Whereas we take a mechanism design approach to our problem, several papers study collective-decision problems in specific bargaining situations. Bergstrom

[6] Cai et al. (2013) show that virtual values can be constructed to describe optimal mechanisms even in settings with multidimensional types, if agent-specific transfers can be used. Our analysis shows that the appropriate notion of a virtual value/cost is substantially simpler in the context of multidimensional IR constraints than in their setting with multidimensional IC constraints.

[7] Also see Gershkov et al. (2017), who further study optimal voting mechanisms for a class of environments with more than two social outcomes.

(1978) studies a setting in which each seller of a commonly owned good names a price to sell their share, and he shows that the likelihood of the good being sold approaches zero. Che (2002) studies how the ability to bargain jointly affects a group's bargaining position. The model takes a hybrid approach in which a group cannot commit to which offers to accept but can commit to a mechanism that specifies how the surplus is divided once an offer is accepted. Grossman and Hart (1980) show that takeover of a firm by a buyer might not be profitable when the buyer offers shareholders a uniform price per share even if the takeover increases efficiency. Oliveros and Iaryczower (2022) study coalition formation when a principal bargains sequentially with a group of agents. Naturally, in many of these collective-decision bargaining games, some form of the holdout problem appears. Instead, in our setting, holdout is implicit and is reflected in the constraint that all sellers must be willing to participate in the mechanism.

2. Model

We study the problem of a buyer who wishes to buy one good, such as a plot of land, from a group of sellers who each own some share of the good. We denote the finite set of agents (the sellers) $N = \{1, \ldots, N\}$ and assume $N \geqslant 2$. Each agent i owns a fraction $\sigma_i \in (0,1)$ of the good, where $\sum_{i \in N} \sigma_i = 1$. The buyer receives a benefit b from purchasing the good. Each agent i's cost of selling his own share of the good (or, equivalently, his valuation for keeping it) is $\sigma_i \boldsymbol{\theta}_i$, where $\boldsymbol{\theta}_i$ denotes the agent's cost per unit of the good.

An outcome of our contracting environment consists of (i) the probability $\mathtt{x} \in [0,1]$ with which the good is sold to the buyer and (ii) the (signed) transfer $\mathtt{m} \in \mathbb{R}$ paid by the buyer to the group of sellers. This transfer is divided among the agents proportionally to their shares, so each agent i receives a payment of $\sigma_i \mathtt{m}$. The assumption that each agent is paid proportionally to his share is motivated by our application to land acquisition, in which a buyer is often required to offer identical terms to sellers ex post. The buyer must treat the agents identically and cannot offer different prices (per unit) to different agents.

The buyer's payoff for outcome $(\mathtt{x}, \mathtt{m})$ is $b\mathtt{x} - \mathtt{m}$. The payoff of each agent i for this outcome is the amount of money he receives minus his cost for his share if the good is sold, $\sigma_i \mathtt{m} - \sigma_i \boldsymbol{\theta}_i \mathtt{x}$. Since σ_i is a positive constant for each i, we can rescale each agent's payoff to be $\mathtt{m} - \mathtt{x}\boldsymbol{\theta}_i$. Such rescaling leaves the agents' incentives unchanged. We henceforth write agent i's payoff as $\mathtt{m} - \mathtt{x}\boldsymbol{\theta}_i$ and refer to $\boldsymbol{\theta}_i$ as agent i's cost.

Let us now describe our informational assumptions. The benefit b is publicly known. Each agent privately knows his own valuation. We assume that the N random variables $\{\boldsymbol{\theta}_i\}_{i \in N}$ are independent and each takes values in the compact interval $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$; denote the cumulative distribution function of $\boldsymbol{\theta}_i$ by F_i.[8] All parties know these distributions.

We make the following regularity assumption for each $i \in N$: the cumulative distribution function F_i admits a density f_i which is continuous and strictly positive, and the *virtual cost* $\varphi_i : \Theta_i \to \mathbb{R}$ given by $\varphi_i(\theta_i) := \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)}$ is strictly increasing. Working directly with an agent's virtual cost $\boldsymbol{\varphi}_i := \varphi_i(\boldsymbol{\theta}_i)$, an atomlessly distributed random variable with convex support, is often convenient. To avoid trivialities, we assume every agent i has $\underline{\theta}_i < b < \varphi_i(\bar{\theta}_i)$.[9]

2.1. Mechanisms

An allocation rule is a measurable function $x : \Theta \to [0, 1]$; let $\mathcal{X}$ denote the set of all allocation rules. A (collective) transfer rule is a bounded measurable function $m : \Theta \to \mathbb{R}$. A (direct) mechanism is a pair (x, m) consisting of an allocation rule and a transfer rule. For any reported type profile θ, the buyer transfers $m(\theta)$ to the group, and $x(\theta)$ is the probability with which she acquires the good.[10]

Say a mechanism (x, m) is **incentive compatible (IC)** if

$$\theta_i \in \operatorname{argmax}_{\hat{\theta}_i \in \Theta_i} \mathbb{E}\left[m(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) - \theta_i x(\hat{\theta}_i, \boldsymbol{\theta}_{-i})\right], \ \forall i \in N, \ \forall \theta_i \in \Theta_i, \tag{IC}$$

that is, report $\hat{\theta}_i = \theta_i$ maximizes the expected payoff of type θ_i of agent i over all possible reports in Θ_i, taking the expectation over other agents' types $\boldsymbol{\theta}_{-i}$. Say the mechanism is **individually rational (IR)** if

$$\mathbb{E}\left[m(\theta_i, \boldsymbol{\theta}_{-i}) - \theta_i x(\theta_i, \boldsymbol{\theta}_{-i})\right] \geqslant 0, \ \forall i \in N, \ \forall \theta_i \in \Theta_i, \tag{IR}$$

that is, the expected payoff of type θ_i of agent i when reporting truthfully, taking the expectation over type profiles of other agents, is nonnegative. An IC and IR mechanism (x, m) generates a buyer profit of

$$\Pi(x, m) := \mathbb{E}\left[bx(\boldsymbol{\theta}) - m(\boldsymbol{\theta})\right].$$

An **optimal mechanism** is an IC and IR mechanism that generates weakly higher

[8] We use the following standard notation throughout. The set of type profiles is $\Theta := \prod_{j \in N} \Theta_j$, and $\Theta_{-i} := \prod_{j \in N \setminus \{i\}} \Theta_j$ for $i \in N$. We also sometimes use a measure and its cumulative distribution function interchangeably, and we use F and F_{-i} to refer to associated product measures on Θ and Θ_{-i}, respectively. Throughout the paper, we use the boldface notation $\boldsymbol{\theta}, \boldsymbol{\theta}_i$, etc. to refer to these random variables, and use the notation θ_i to refer to an element of Θ_i (that is, a potential realization of $\boldsymbol{\theta}_i$).

[9] This assumption reduces casework but is not important for analyzing our model. For example, without it, Theorem 1 would still hold as stated, except that when the essentially unique allocation rule specifies never trading or always trading, the weights can be non-unique.

[10] Because the payoffs are linear in the transfer, we assume without loss that in a direct mechanism the transfer is a deterministic function of the reported type profile.

buyer profit than any other IC and IR mechanism. An **optimal allocation rule** is any allocation rule x such that (x, m) is an optimal mechanism for some m.

2.2. Alternative interpretations of our model

Before moving on to our analysis, we discuss some alternative interpretations of our model.

Recall that after normalization, we have a setting in which the buyer's payoff for outcome $(\mathtt{x}, \mathtt{m})$ is $b\mathtt{x} - \mathtt{m}$ and each agent i's payoff is $\mathtt{m} - \boldsymbol{\theta}_i\mathtt{x}$. We can interpret this setting as one in which the agents sell a good they collectively own in exchange for *public funds* (instead of some money that is divided between them). The money $\mathtt{m}$ appears in every agent's payoff function. For example, the agents might be a committee of decision-makers in an organization, such as the high-level executives at a firm, who decide whether they should sell an asset owned by the organization. In this interpretation, an agent's type $\boldsymbol{\theta}_i$ specifies his marginal rate of substitution between the organization retaining the good and the organization's use of additional funds.

We do not make assumptions about the signs of b or the values $\boldsymbol{\theta}_i$ might take. In particular, we allow them to be negative. In that case, after relabeling the variables appropriately, the problem becomes one of finding optimal mechanisms for a single *seller* who wants to sell a good to a group of buyers. If sold, the good is publicly available to all members of the group. Depending on whether the agents in the group pay for the good with private money or public funds, two interpretations are again available. The first interpretation entails private transfers with a fixed cost-sharing rule. Here, each agent i is responsible for paying a fixed fraction σ_i of the transfer to the seller. So if the good is sold with probability $\mathtt{x}$ and the group pays $\mathtt{m}$ to the seller, then agent i's payoff is $\boldsymbol{v}_i\mathtt{x} - \sigma_i\mathtt{m}$, where $\boldsymbol{v}_i$ denotes agent i's benefit if the good is acquired by the group. For example, the group might be a condo association in which each member pays for a public service proportionally to the size of their unit, or it might be a cartel in which each firm pays proportionally to its market share. Now define $\boldsymbol{\theta}_i := \frac{1}{\sigma_i}\boldsymbol{v}_i$, which allows us to write agent i's payoff as $\sigma_i\boldsymbol{\theta}_i\mathtt{x} - \sigma_i\mathtt{m}$, which can then be normalized to $\boldsymbol{\theta}_i\mathtt{x} - \mathtt{m}$. The second interpretation has the group paying for the product with public funds. Here, if the good is sold with probability $\mathtt{x}$ and the group pays $\mathtt{m}$ to the seller from its collective funds, then agent i's payoff is $\boldsymbol{\theta}_i\mathtt{x} - \mathtt{m}$. Agent i's type $\boldsymbol{\theta}_i$ again denotes his marginal rate of substitution between the public good and the organization's alternative use of its funds.

3. Characterizing the optimal mechanism

In this section, we fully characterize optimal mechanisms. First, we describe which allocation rules are implementable and solve for the buyer's optimal profit from implementing such an allocation rule; doing so requires a reduced-form implementation result for transfers, characterizing exactly which profiles of interim transfer rules can be implemented with some collective transfer rule. Then, the main result of this section establishes that a unique optimal allocation rule exists, and shows it can be described as a weighted allocation rule with (uniquely determined) weights that we explicitly characterize.

We begin by introducing some convenient notation and terminology for standard objects. Just as in the auction setting, the Bayesian incentive properties of our design environment are convenient to discuss in terms of each agent's interim (i.e., conditioning only on his own type) outcomes.

Definition 1: *Fix any agent $i \in N$. Given an allocation rule x, define the* ***interim allocation rule*** *to be $X_i^x : \Theta_i \to \mathbb{R}$ given by $X_i^x(\theta_i) := \mathbb{E}[x(\theta_i, \boldsymbol{\theta}_{-i})]$. Similarly, given a transfer rule m, define the* ***interim transfer rule*** *to be $M_i^m : \Theta_i \to [0,1]$ given by $M_i^m(\theta_i) := \mathbb{E}[m(\theta_i, \boldsymbol{\theta}_{-i})]$.*

Now, say an allocation rule x is ***interim monotone*** *if X_i^x is weakly decreasing for every $i \in N$. We say that an allocation rule x is* ***implementable*** *if a transfer rule m exists such that (x, m) is IC.*

To characterize optimal mechanisms, we first need to understand which allocation rules are implementable. Classic results (Myerson, 1981; Myerson and Satterthwaite, 1983) would imply that interim monotonicity would fully characterize implementability if the buyer could freely choose the interim transfer rule that each agent faces. However, our buyer is constrained in that different agents' interim transfers must be derived from a common ex post transfer rule. Nevertheless, Lemma 1 below shows that the exact same characterization applies despite the added constraint on the transfers. Using this characterization, we also obtain the maximum buyer profit compatible with implementing an allocation rule x.

Lemma 1: *Let x be some allocation rule.*

(i) Mechanism (x, m) is IC and IR for some transfer rule m if and only if x is interim monotone.

(ii) If some transfer rule m exists such that mechanism (x, m) is IC and IR, then a maximally profitable such mechanism exists, with resulting profit

$$\min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right].$$

Part (i) of the lemma combines a standard observation with a novel one. The standard observation is that a given mechanism (x, m) is IC for seller i if and only if X_i is weakly decreasing and some constant $\underline{U}_i$ exists such that the payment identity,

$$M_i(\theta_i) = \underline{U}_i + X_i(\theta_i)\theta_i + \int_{\theta_i}^{\bar{\theta}_i} X_i(\tilde{\theta}_i)\, \mathrm{d}\tilde{\theta}_i \text{ for every } \theta_i, \tag{\$}$$

holds. So the allocation rule x and the constants $\underline{U}_1, \ldots, \underline{U}_N$ pin down the *interim* transfer rules. The novel observation for establishing part (i) is that a profile of interim transfer rules $(M_i)_{i \in N}$ can be implemented via a common ex-post transfer rule m if and only if $\mathbb{E}[M_i(\boldsymbol{\theta}_i)]$ coincide for all i. This condition is obviously necessary, given iterated expectations. Perhaps surprisingly, the condition is also sufficient, and sufficiency has a one-line proof: if $\bar{m}$ is the common expected transfer, then the ex-post transfer rule $m(\theta) := -(N-1)\bar{m} + \sum_{i \in N} M_i(\theta_i)$ generates the desired interim transfer rules.[11]

To establish part (ii), we use the payment identity (\$) to write the buyer's expected payoff as

$$\mathbb{E}\left[bx(\boldsymbol{\theta}) - m(\boldsymbol{\theta})\right] = \mathbb{E}\left[bx(\boldsymbol{\theta})\right] - \mathbb{E}\left[M_i(\boldsymbol{\theta}_i)\right] = \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] - \underline{U}_i$$

for each seller i. Importantly, because the interim transfers are identical on average, choosing $\underline{U}_i$ for any seller i pins down the entire profile of constants $(\underline{U}_i)_i \in \mathbb{R}^N$. Analogously to how an optimal auction would optimize the transfer rule by setting each agent's IR constraint to be binding, our remaining constant is optimized by requiring that *some* agent's IR constraint binds (and the others' constraints are satisfied). Since $\mathbb{E}[M_i(\boldsymbol{\theta}_i)] = \mathbb{E}[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i] + \underline{U}_i$ coincide for all $i \in N$, an agent i whose IR constraint binds is the one with the highest expected virtual cost of trade $\mathbb{E}\left[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i\right]$. Therefore, the buyer's optimal payoff for a given allocation rule x is

$$\min_{i \in N} \; \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right].$$

With Lemma 1 in hand, our buyer's problem can be recast directly as an optimization over allocation rules. Formally, the buyer's optimization over allocation

[11] Gopalan et al. (2018) show that a slight variant of this problem is computationally intractable. In particular, it is computationally hard to decide whether a given profile of interim transfer rules can be implemented via a common ex post transfer *that is constrained to belong to some bounded interval.* Our construction—which settles the question of implementability absent such a constraint—resembles previous constructions in the literature that convert ex-ante budget-balanced mechanisms into ex post budget-balanced mechanisms while preserving the players' interim transfer rules (e.g., Makowski and Mezzetti, 1994; d'Aspremont et al., 2004; Che and Kim, 2006; Börgers and Norman, 2009).

rules is

$$\max_{x \in \mathcal{X}} \left\{ \min_{i \in N} \ \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] \right\} \tag{BP}$$
$$\text{s.t.} \quad x \text{ is interim monotone.}$$

Our main result is a complete characterization of the solution to the program (BP). To this end, we define a class of allocation rules that play a special role in our analysis and results.

Definition 2: *Given $\omega \in \Delta N$, the **ω-allocation rule** is the allocation rule $x_\omega := \mathbb{1}_{\omega \cdot \varphi \leqslant b}$. Say $\omega \in \Delta N$ is **optimal** if the ω-allocation rule is optimal. Say an allocation rule is a **weighted allocation rule** if it is a ω-allocation rule for some $\omega \in \Delta N$.*

We now state our main characterization theorem.

Theorem 1 (Optimal allocation)**:**
A weighted allocation rule is essentially uniquely optimal.[12] *The unique optimal-weight vector ω is characterized by either of the following two equivalent conditions:*

(i) $\omega \in \operatorname{argmin}_{\tilde{\omega} \in \Delta N} \mathbb{E}[(b - \tilde{\omega} \cdot \boldsymbol{\varphi})_+]$.

(ii) $\operatorname{supp}(\omega) \subseteq \operatorname{argmax}_{i \in N} \mathbb{E}\left[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b\right]$.

Moreover, if $b < \bar{\theta}_j$ for at least two $j \in N$, then every $i \in N$ has $\omega_i < 1$.

Theorem 1 says that trade outcomes in optimal mechanisms are given by weighted allocation rules. Notice that a weighted allocation rule is deterministic, so the theorem implies the buyer does not benefit from randomization. Note also that the weights are fixed and do not depend on reports. A higher weight for an agent means the mechanism is in a sense more responsive to that agent's private information. So by comparing agents' weights, we can understand which agents exert greater influence over the outcomes of the mechanism, a topic we will revisit in the next section.

Theorem 1 also characterizes the optimal weights with two equivalent conditions. Condition (i) above is useful for computing the optimal weights numerically and analytically. The function $\omega \mapsto \mathbb{E}[(b - \omega \cdot \boldsymbol{\varphi})_+]$ is convex, and so determining optimal weights corresponds to minimizing a convex objective over a compact convex set. Moreover, in certain cases, as in Example 1 presented later, we can even compute the optimal weights ω analytically using the first-order conditions of the convex optimization problem. Condition (ii) of the theorem facilitates verification of optimality: once a candidate for optimal weights is chosen, one can verify opti-

[12] By "essentially uniquely," we mean any alternative optimal allocation rule generates the same trade decision almost surely.

mality by checking that any agent who has a positive weight is a minimizer of the conditional expected virtual cost term. This condition reflects the fact that every agent who influences the trade outcome should have a binding IR constraint.

Specializing to the case of a single seller, Theorem 1 confirms the classic characterization of optimal mechanisms. In this case, trade happens whenever the benefit to the buyer exceeds the lone seller's virtual cost. The "weight" for this case is trivial, assigning all influence the seller's private information. The optimal allocation is deterministic (trade occurs whenever the agent's type is below a cutoff type at which the benefit is equal to virtual cost), just as in Theorem 1.

The proof of Theorem 1 studies a relaxed program (RBP) in which the interim-monotonicity constraint is ignored. To solve the relaxed program, we consider an auxiliary two-player zero-sum game in which the Maximizer chooses an allocation rule x, the Minimizer chooses an agent i whose IR constraint must bind, and so the Maximizer's objective is $\mathbb{E}\left[x(\boldsymbol{\theta})(b-\boldsymbol{\varphi}_i)\right]$. Observe that an allocation rule solves (RBP) if and only if it is a cautious optimum for the Maximizer in the auxiliary game—that is, a maximin strategy. Moreover, standard results for zero-sum games imply a maximin strategy is a Nash equilibrium strategy for the Maximizer, and vice versa, as long as some Nash equilibrium exists. Hence, we turn to characterizing Nash equilibria of the auxiliary game.

We first show that if the Minimizer is allowed to choose a mixture, some Nash equilibrium of this auxiliary game exists by an appropriate minimax theorem, and every mixed strategy ω for the Minimizer exhibits a unique (up to almost-sure equivalence) best response for the Maximizer. Indeed, x_ω is the essentially unique maximizer of

$$x \mapsto \sum_i \omega_i \mathbb{E}\left[x(\boldsymbol{\theta})(b-\boldsymbol{\varphi}_i)\right] = \mathbb{E}\left[x(\boldsymbol{\theta})(b-\omega\cdot\boldsymbol{\varphi})\right]$$

because it sets $x(\boldsymbol{\theta}) \in [0,1]$ to maximize the integrand $x(\boldsymbol{\theta})(b-\omega\cdot\boldsymbol{\varphi})$ in every realized state. Then, because the set of Nash equilibria of a two-player zero-sum game exhibits a product structure, it follows that an essentially unique allocation rule can be an optimal strategy for the Maximizer of the auxiliary game, and that it takes the form x_ω for any Nash equilibrium choice ω of the Minimizer. The pair of conditions characterizing such ω's are standard to zero-sum games: the mixed strategy ω is a cautious optimum for the Minimizer (condition (i)) if and only if it is a best response to some Maximizer's best response to ω (condition (ii), once the Maximizer's best response to ω is substituted in). Now, observe that the essentially unique Nash equilibrium strategy for the Maximizer is actually interim monotone: because virtual costs are increasing, a cutoff rule for the ω-weighted virtual cost is monotone and hence interim monotone. The result is a characterization of the unique optimal allocation rule, solving not only (RBP) but also (BP). Then, because our assumption that $\underline{\theta}_i < b < \varphi_i(\bar{\theta}_i)$ (for each i) implies every weighted allocation rule has an interior probability of trade, a geometric argument converts uniqueness of the allocation rule into uniqueness of agents' weights. Finally, to verify the last sentence of the theorem, we note that any agents $i \neq j$ have $\mathbb{E}[\boldsymbol{\varphi}_i \mid \boldsymbol{\varphi}_i \leqslant b] \leqslant b$ and $\mathbb{E}[\boldsymbol{\varphi}_j \mid \boldsymbol{\varphi}_i \leqslant b] = \mathbb{E}[\boldsymbol{\varphi}_j] = \bar{\theta}_j$, so that putting all weight on agent i would violate condition (ii) if $\bar{\theta}_j > b$.

To conclude the section, let us specialize our setting to a parametric example. We use the example to illustrate how we can use condition (i) of Theorem 1 to identify the optimal weights analytically. We then give an indirect implementation of the optimal mechanism.

Example 1: *Suppose that there are two sellers. Seller i has a power distribution $F_i(\theta_i) = \theta_i^{\alpha_i}$ over $\theta_i \in [0, 1]$ for some power $\alpha_i > 0$.*[13] *Assume that $b \leqslant \min\{\frac{1+\alpha_1}{\alpha_1+\alpha_2}, \frac{1+\alpha_2}{\alpha_1+\alpha_2}\}$—which for instance holds if $\alpha_1, \alpha_2, b \leqslant 1$. Then, as we show in the appendix, the optimal weight vector is:*

$$\omega^* := \left(\frac{\alpha_1}{\alpha_1 + \alpha_2}, \quad \frac{\alpha_2}{\alpha_1 + \alpha_2} \right).$$

A convenient feature of this example is that virtual costs are linear in costs,

$$\varphi_i(\theta_i) = \left(1 + \frac{1}{\alpha_i}\right) \theta_i.$$

Combining this observation with the optimal weights we have computed, the essentially unique optimal allocation rule results in trade if and only if

$$b \geqslant \omega^* \cdot \boldsymbol{\varphi} = \frac{\alpha_1}{\alpha_1 + \alpha_2} \boldsymbol{\varphi}_1 + \frac{\alpha_2}{\alpha_1 + \alpha_2} \boldsymbol{\varphi}_2 = \frac{1 + \alpha_1}{\alpha_1 + \alpha_2} \boldsymbol{\theta}_1 + \frac{1 + \alpha_2}{\alpha_1 + \alpha_2} \boldsymbol{\theta}_2.$$

So trade occurs whenever the benefit is at least a certain positive linear combination of the sellers' costs. This allocation rule, combined with any transfer rule that satisfies the payment identity ($) with $\underline{U}_1 = \underline{U}_2 = 0$, forms an optimal mechanism.

In addition, we show that the following indirect mechanism is optimal:

- *Both sellers simultaneously send bids, $s_i \in \mathbb{R}_+$.*
- *Trade occurs if and only if $b \geqslant \tau_1 s_1 + \tau_2 s_2$, where $\tau_i := \frac{1+\alpha_{-i}}{(\alpha_1+\alpha_2)\alpha_{-i}}$ for each i.*
- *The price that the buyer pays is $s_1 + s_2 + \kappa b$ if the good is sold, and zero*

[13] When $\alpha_i \neq 1$, the derivative of F_i at $\underline{\theta}_i = 0$ is either zero or infinite, violating our assumption of a continuous and strictly positive density. However, our main results (in particular, Theorems 1 and 2 and their supporting analysis) apply to the more general version of our model in which the density is only assumed to be continuous, finite, and strictly positive on $(\underline{\theta}_i, \bar{\theta}_i)$. In particular, assuming the density f is positive over $(\underline{\theta}_i, \bar{\theta}_i)$, we can define $\varphi_i(\theta_i) := \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)}$ over $(\underline{\theta}_i, \bar{\theta}_i)$, and extend it continuously to the endpoints.

otherwise, where $\kappa = \frac{\alpha_1+\alpha_2}{(1+\alpha_1)(1+\alpha_2)}(1-\alpha_1\alpha_2)$.

This indirect mechanism has a "name-your-price" structure. Each seller submits a bid. It is useful to think of the bid as the price the buyer has to pay in exchange for that seller's consent. The buyer pays the sum of the bids (plus a constant) if and only if trade occurs. Trade occurs when the benefit is higher than some linear combination of the bids.

When submitting a bid, each seller faces a trade-off like in a first-price auction. Increasing the bid means the agent is paid more if *trade occurs. But increasing the bid also lowers the probability of trade. We show that the game has an equilibrium with linear strategies in which type* θ_i *bids* $\frac{(1+\alpha_i)\alpha_{-i}}{1+\alpha_{-i}}\theta_i$*, and in this equilibrium the buyer obtains her maximum possible payoff.*

As mentioned earlier, determining the optimal weights using condition (i) in Theorem 1 entails solving a convex optimization program. We show in the appendix that ω^* is a local minimum of this program, and hence is optimal. Toward showing the indirect mechanism is optimal, we first establish that the given strategy profile generates the above-described allocation rule and interim allocation rules. Thus, the indirect mechanism and strategy profile constitute an optimal mechanism *if* the strategy profile is an equilibrium. Moreover, because the induced allocation rule is interim monotone and the payment identity ($) holds, it follows that no type of either agent has a profitable deviation to submit another type's bid. Finally, notice that any off-path bid—a bid not submitted by any type on-path—is outcome equivalent to bidding the highest type's bid (generating no trade and zero transfer). Therefore, no type has any profitable deviation—that is, the given strategy profile is indeed an equilibrium.

Let us highlight two important features of the above example. First, notice that $\omega_1^* > \omega_2^*$ whenever $\alpha_1 > \alpha_2$: the seller with a higher α_i receives a higher weight in the optimal mechanism, and in this sense has greater influence over the outcomes of the mechanism. In the bidding-game implementation, this influence ranking is reflected in $\tau_1 > \tau_2$. Section 4 studies the role of asymmetry more broadly, asking how ex-ante heterogeneity in sellers' characteristics leads to asymmetric treatment by the mechanism beyond this parametric example. Using the characterization of optimal weights in Theorem 1, we show that the optimal mechanism assigns a higher weight to agents who have higher costs ex ante. Theorem 2 formalizes the appropriate sense in which $\alpha_1 \geqslant \alpha_2$ corresponds to seller 1 having higher costs ex ante, and shows that it generates a ranking of weights more generally.

Second, the terms of trade arose in the example from a complex pricing mechanism—complex in the sense that bidding behavior altered the terms of trade, not just whether trade occurred. In particular, the price at which the trade occurs depends on the type profile, unlike the case with $N = 1$ where the optimal mechanism can be implemented with a posted price. In Section 5, we show that this complex pricing aspect of this mechanism is a feature of every optimal mechanism in our setting.

4. The role of agent heterogeneity

In light of our leading application—sale of a large plot of land with dispersed ownership to an industrialist—it is natural to explore how the optimal mechanism treats (ex ante) heterogeneity between agents. Because the optimal mechanism uses a weighted-average allocation rule, this question amounts to understanding how agents' endogenous weights differ.

Specifically, we seek conditions on primitives under which we can rank ω_i and ω_j for two agents i and j. The main result of this section, Theorem 2, provides an interpretable condition under which a ranking of agents' virtual cost distributions implies a ranking on the weights in the optimal mechanism. To state this result, we use the following distributional-ranking definition.

Definition 3: *Given two real random variables* $\mathbf{v}$ *and* $\mathbf{w}$ *with respective cumulative distribution functions given by* G *and* H*,* $\mathbf{v}$ ***is larger than*** $\mathbf{w}$ ***in the reversed hazard-rate order****, denoted by* $\mathbf{v} \succcurlyeq_{\text{rh}} \mathbf{w}$*, if* $\inf[\text{supp}(\mathbf{w})] \leqslant \inf[\text{supp}(\mathbf{v})]$ *and* $\frac{G}{H}$ *is weakly increasing on* $(\inf[\text{supp}(\mathbf{w})], \infty)$*.*

The above distributional ranking is a useful strengthening of first-order stochastic dominance. Intuitively, the ranking requires that the conditional distributions, when conditioned on lying below any common threshold, are stochastically ranked. This ranking condition has been fruitful in past work in mechanism design. Specifically, in the literature on asymmetric auctions (e.g., Maskin and Riley, 2000; Kirkegaard, 2012), ranking bidders' value distributions via the reversed hazard-rate order has enabled the ranking of equilibrium bidding behavior, which in turn has been used to provide revenue rankings for alternative auction formats. In our setting, as the following theorem shows, a reversed hazard-rate order on agents' *virtual cost* distributions is relevant in designing optimal mechanisms.

Theorem 2 (Ranking allocation weights)**:** *If* $\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \boldsymbol{\varphi}_j + \beta$ *for some* $\beta \geqslant 0$*, then the optimal vector of allocation weights* ω *satisfies* $\omega_i \geqslant \omega_j$*. Moreover,* $\omega_i > \omega_j$ *whenever* $\beta > 0$ *and* $\omega_j > 0$*.*

Theorem 2 follows from more general results (which further provide quantitative bounds on *how* asymmetric the weights are) that we prove in the appendix.[14] The core of the theorem's proof is a result from the theory of stochastic orders that

[14] More specifically, if $\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \alpha\boldsymbol{\varphi}_j + \beta$, where $0 < \alpha \leqslant 1$ and β exceeds a certain bound, then $\alpha\omega_i \geqslant \omega_j$; and we prove a corresponding strict version of the same.

converts a reversed hazard-rate ranking on random variables into a second-order stochastic-dominance ranking of their weighted averages as the weights are made more assortative. More specifically, we work with the convex program given in condition (i) of Theorem 1, and note that its loss function can be written as

$$-\mathbb{E}\left[h(\omega_i \boldsymbol{\varphi}_i + \omega_j \boldsymbol{\varphi}_j)\right]$$

for some increasing and concave function h that depends on $(\omega_k)_{k \neq i,j}$. Suppose $\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \boldsymbol{\varphi}_j + \beta$ for some $\beta \geqslant 0$, and consider any weight vector ω with $\omega_i < \omega_j$. We can then define an alternate weight vector $\tilde{\omega}$ by swapping the i and j coordinates of ω. Because $\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \boldsymbol{\varphi}_j + \beta$ and $\omega_i < \omega_j$, a textbook stochastic ranking result tells us $\omega_i \boldsymbol{\varphi}_i + \omega_j(\boldsymbol{\varphi}_j + \beta)$ is below $\omega_j \boldsymbol{\varphi}_i + \omega_i(\boldsymbol{\varphi}_j + \beta)$ in the sense of second-order stochastic dominance. But then,

$$\mathbb{E}\left[h\left(\omega_i \boldsymbol{\varphi}_i + \omega_j \boldsymbol{\varphi}_j\right)\right] \leqslant \mathbb{E}\left[h\left(\omega_j \boldsymbol{\varphi}_i + \omega_i \boldsymbol{\varphi}_j - (\omega_j - \omega_i)\beta\right)\right] \leqslant \mathbb{E}\left[h\left(\omega_j \boldsymbol{\varphi}_i + \omega_i \boldsymbol{\varphi}_j\right)\right],$$

so that $\tilde{\omega}$ performs at least as well as ω in the convex program. But Theorem 1 then tells us that ω cannot be the unique optimal weight vector.

When types are drawn from power distributions, virtual costs can be ranked in the reversed hazard-rate order, and so Theorem 2 gives a ranking of the weights that matches our closed-form calculations in Example 1. In particular, consider two agents i and j with distributions given by $F_i(\theta_i) = \theta_i^{\alpha_i}$ and $F_j(\theta_j) = \theta_j^{\alpha_j}$ for $\theta_i, \theta_j \in [0,1]$, where $\alpha_i, \alpha_j > 0$. If $\alpha_i \geqslant \alpha_j$, then $\boldsymbol{\varphi}_i$ is higher than $\boldsymbol{\varphi}_j$ in the reversed hazard-rate order, and so $\omega_i \geqslant \omega_j$ by the theorem, regardless of the distributions of other agents.[15]

Let us now revisit our land-acquisition interpretation. The principal, an industrialist, wishes to buy a large plot of land whose ownership is dispersed across N individuals, with agent i owning share σ_i of the land. Each agent's valuation per unit of land is $\boldsymbol{\theta}_i$, and his utility is $\sigma_i \mathtt{m} - \sigma_i \boldsymbol{\theta}_i \mathtt{x}$. Because σ_i is a positive scalar multiplying $\mathtt{m} - \boldsymbol{\theta}_i \mathtt{x}$, it is strategically irrelevant. Therefore, if the agents' virtual cost distributions are ranked according to the reversed hazard-rate order, then a ranking of the weights follows. Land shares per se play no role in determining agents' optimal weights. For instance, if two agents have the same virtual cost distributions, then the optimal mechanism will weigh them equally however asymmetric their land shares are.

[15] As pointed out in Example 1, any $\theta_i \in [0,1]$ has $\varphi_i(\theta_i) = \frac{\alpha_i+1}{\alpha_i}\theta_i$, implying $\mathbb{P}[\boldsymbol{\varphi}_i \leqslant x] = \left(\frac{\alpha_i x}{\alpha_i+1}\right)^{\alpha_i}$ for $x \in \left[0, \frac{\alpha_i+1}{\alpha_i}\right]$—and analogously for j. Because $\frac{\alpha_i+1}{\alpha_i} \geqslant \frac{\alpha_j+1}{\alpha_j}$ and any x in the common support has

$$\frac{\mathbb{P}[\boldsymbol{\varphi}_i \leqslant x]}{\mathbb{P}[\boldsymbol{\varphi}_j \leqslant x]} \text{ proportional to } x^{\alpha_i - \alpha_j},$$

it follows that $\boldsymbol{\varphi}_i$ is larger than $\boldsymbol{\varphi}_j$ in the reversed hazard-rate order.

But could the amount of landholding be systematically related to the cost distribution? For example, consider two sellers with landholdings $\sigma_1 < \sigma_2$, and assume that the shares are sufficiently asymmetric. Then it is conceivable that the agent with a larger landholding may have uses of land that generate higher value (thus, a higher cost of trade) in per-unit terms. For example, an agent with a larger piece of land might install a manufacturing plant. The smaller landowner cannot do the same because of the associated fixed costs and minimum-size constraints. This difference in how they use their plots can lead to $\boldsymbol{\varphi}_2 \geqslant_{\text{rh}} \boldsymbol{\varphi}_1$; that is, the agent with a larger landholding may have higher productivity (and therefore cost of trade) per unit of land. As Theorem 2 says, the optimal mechanism assigns agent 2, the more productive agent, a higher weight.

Another compelling story could, however, apply to situations in which all the landowners have the same land use, say agriculture, and they differ only in the sizes of plots they own (in addition to idiosyncratic shocks). A negative relationship between the size of land and productivity is well documented (e.g., Banerjee, 1999; Berry et al., 1979). In fact, the magnitude of this difference in productivity is often sizable. As Banerjee (1999) says:[16]

> In Punjab, Pakistan, productivity on the largest farms (as measured by value added per unit of land) is less than 40 percent that on the second smallest size group, while in Muda, Malaysia, productivity on the largest farms is just two-thirds that on the second smallest size farms.

In such contexts, in which the agents with smaller landholdings are more productive (in per-unit terms), we could have $\boldsymbol{\varphi}_1 \geqslant_{\text{rh}} \boldsymbol{\varphi}_2$ and therefore $\omega_1 \geqslant \omega_2$ per Theorem 2. That is, the optimal mechanism would favor the agents with smaller landholdings.

In summary, a general qualitative feature emerges from the above two situations: the optimal mechanism favors the more productive agents, who are less ex ante inclined to part with their land. Given a systematic positive or negative relationship between agents' productivity and their landholdings, this observation further enables us to understand which agents the optimal mechanism favors.

5. Posted-price mechanisms

In some mechanism design problems—for example, selling a single indivisible good to a single agent—the optimal mechanism is a take-it-or-leave-it posted-price mechanism (Myerson, 1981; Riley and Zeckhauser, 1983). Beyond the single-

[16] One reason Banerjee (1999) offers for this negative relationship is decreasing returns to scale arising from incentive costs. Smaller plots tend to be managed by families, while larger ones require significant external labor.

agent setting, there are environments in which such pricing mechanisms remain approximately optimal (Chawla et al., 2010; Chawla et al., 2015; Hart and Nisan, 2017; Babaioff et al., 2020). Especially in our setting—in which any agent can unilaterally veto the mechanism and all the agents must pay a common price—a natural conjecture is that posted-price mechanisms remain optimal. The purpose of this section is to establish that this conjecture is false. Of course, before we can do so, we must first define a posted-price mechanism for our setting.

In the one-agent setting, the IC direct mechanisms that correspond to a posted price are those satisfying two properties. First, the transfer is directly proportional to the allocation probability. And second, the allocation probability is 1 for types above the price and 0 for those below it. The first condition—which we can interpret as a restriction that money never changes hands if the good is not sold and that the price at which trade occurs is constant when it does—generalizes immediately. But the second condition—which we can interpret as stating that the agent freely decides whether to execute a trade—does not immediately generalize to the multi-agent setting. Who decides whether trade occurs? Once the buyer announces a price for the good, a complex negotiation process could ensue between the agents deciding whether to sell. Might eventual trade outcomes arise from some mixed-strategy equilibrium of the resulting bargaining game between the agents?

In light of these difficulties, we define a collective posted price rather permissively, only incorporating the first of the two conditions mentioned in the previous paragraph. We also introduce a specific, interpretable pricing mechanism that will be important for our results.

Definition 4: *A mechanism (x, m) is a* ***collective posted-price mechanism*** *if some $p \in \mathbb{R}_+$ exists such that $m = px$. It is a* ***unanimous posted-price mechanism*** *if it is a collective posted-price mechanism with price p such that $x(\theta) = \mathbb{1}_{\theta_i \leqslant p\ \forall i \in N}$ for every $\theta \in \Theta$.*

One can envision several examples of collective posted-price mechanisms. For example, the buyer could set a price p and execute a sale if and only if all agents agree to the purchase—a unanimous posted price. Alternatively, the buyer could post a price and select an agent, or even a subset of agents, perhaps randomly, and execute the trade if all the agents in this chosen subset agree to the sale.

Although the space of all collective posted-price mechanisms is rather rich, the next result shows that arguably the simplest example of them is optimal.

Proposition 1 (Optimal posted price is unanimous)**:** *Some unanimous posted-price mechanism is optimal among IC and IR collective posted-price mechanisms.*

To show this result, we begin with an arbitrary collective posted-price mechanism, with a view to showing some unanimous posted price does better. If the price exceeds the benefit of trade b, then the mechanism is not profitable, and so a unanimous posted price slightly below b yields higher buyer profit. Now focus on the case of a price below b. Observe that IR implies an agent's interim allocation is zero whenever the agent's type is above the price. Therefore, trade has zero probability conditional on *any* agent's having a realized valuation above the price. Hence, a unanimous posted price (at the same price level) would generate profitable trade with a higher probability, and so it is more profitable.

Having characterized the optimal form of collective posted-price mechanism, we are poised to answer the question that motivated this subsection: when are collective posted-price mechanisms optimal? The result below establishes that, under a mild nondegeneracy condition, they never are.

PROPOSITION 2 (Posted prices are suboptimal)**:** *If at least two $j \in N$ have $b < \bar{\theta}_j$, then no collective posted-price mechanism is optimal.*

To establish the above result, in light of Proposition 1, it suffices to show the optimal allocation mechanism is not a unanimous posted price. We show the two cannot coincide by examining their interim allocation rules for an agent who has positive weight in the optimal weighted allocation rule. His interim allocation rule is clearly a step function under a unanimous posted-price mechanism. Meanwhile, it cannot be one under the optimal mechanism: his interim probability of trade is nonconstant in his type because the allocation rule puts positive weight on his own virtual cost, and it is continuous in his type because it puts positive weight on the (atomlessly distributed) virtual cost of at least one other agent. So the two allocation rules cannot coincide.

Thus, optimal mechanisms incorporate sellers' private information, smoothly varying the terms of trade with an agent's reported type. Observe that such continuous incentives are reflected in the bidding game of Example 1, in which changing a bid leads to a change in the price conditional on trade.

6. Group versus single-agent mechanisms

In this section, we compare optimal mechanisms in our setting with optimal mechanisms for buying from a *single* agent. We use this comparison to highlight how optimal allocations are distorted and discuss welfare consequences of these distortions.

We start by defining the single-agent benchmark. In this benchmark, a single agent owns a good that has valuation $\mathbf{v} := \sigma \cdot \boldsymbol{\theta}$ for him, where $\sigma = (\sigma_1, \ldots, \sigma_N)$ is a fixed vector of positive weights summing to 1, and the random vector $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_N)$ has each $\boldsymbol{\theta}_i$ drawn independently from F_i. Using the terminology from our land acquisition application, the agent owns a plot of land that is divided into

N parcels, possibly of different sizes. Land parcel i has size σ_i and selling it has per-unit cost $\boldsymbol{\theta}_i$ to the agent.[17] The agent privately knows $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_N$ (and hence privately knows $\mathbf{v}$). The buyer designs a (direct) mechanism in which the single seller reports her type $\boldsymbol{\theta}$, resulting in a probability of trade and a transfer the buyer pays him. We study mechanisms that maximize the buyer's profit subject to single-agent analogues of the IC and IR constraints. Let G denote the cumulative distribution function $\mathbf{v}$, and let g be the continuous and strictly positive density of $\mathbf{v}$. Although we do not make this assumption for our analysis, let us focus our discussion around the regular case in which the associated virtual cost $v + \frac{G(v)}{g(v)}$ is strictly increasing in v.[18]

Notice that, both in the single-agent benchmark and in our group setting, a mechanism stipulates a probability of trade as a function of the random variable $\boldsymbol{\theta}$, and that this random variable has the same distribution in both settings. Also, in both settings, the total monetary value of the good to the seller(s) is $\sigma \cdot \boldsymbol{\theta}$ and has the same distribution in both cases. From the perspective of the buyer, regardless of whether she interacts with a single agent (as described in the previous paragraph) or with the group (as in our main model), she is paying money to buy a good, and she has the same belief about how valuable the good is to the seller(s). Hence, the utilitarian-efficient allocation is the same in either setting: the good is efficiently traded whenever the benefit of doing so is greater than its cost,

$$b \geqslant \sigma \cdot \boldsymbol{\theta}.$$

For any c_0, let the **iso-cost curve for** c_0 be the set of all type profiles that have the same cost c_0, i.e., those $\theta \in \Theta$ that satisfy $c_0 = \sigma \cdot \theta$. Then, efficient trade occurs in the region of the type space that is below the iso-cost curve for b. Panel (a) of Figure 1 illustrates this region for the case of two agents ($N = 2$), in which iso-cost curves are straight lines with slope $-\frac{\sigma_1}{\sigma_2}$. Iso-cost curves that are further in the northeast direction correspond to larger cost levels. In what follows, we compare the efficiency of the buyer's optimal allocation rule across these two models.

[17] These costs might be independent (conditional on observables) if they represent productivity of different parcels of land, and any shocks that affect multiple parcels' productivity are observable to the buyer.

[18] We make this assumption here only to streamline the exposition. Regularity of the distribution of $\mathbf{v}$ would follow if we were to assume each $\frac{f_i}{F_i}$ is nonincreasing on $(\underline{\theta}_i, \bar{\theta}_i)$ for each $i \in N$. Indeed, in this case, the corresponding ratio for $\mathbf{v}_i = \sigma_i \boldsymbol{\theta}_i$ given by

$$v_i \mapsto \frac{f_i\left(\frac{1}{\sigma_i} v_i\right)}{\sigma_i F_i\left(\frac{1}{\sigma_i} v_i\right)}$$

is nonincreasing there too. Iteratively applying Corollary 3.3 from Barlow et al. (1963) then tells us $\frac{g}{G}$ is nonincreasing there, so that the associated virtual cost is strictly increasing—that is, $\mathbf{v}$ has a regular distribution.

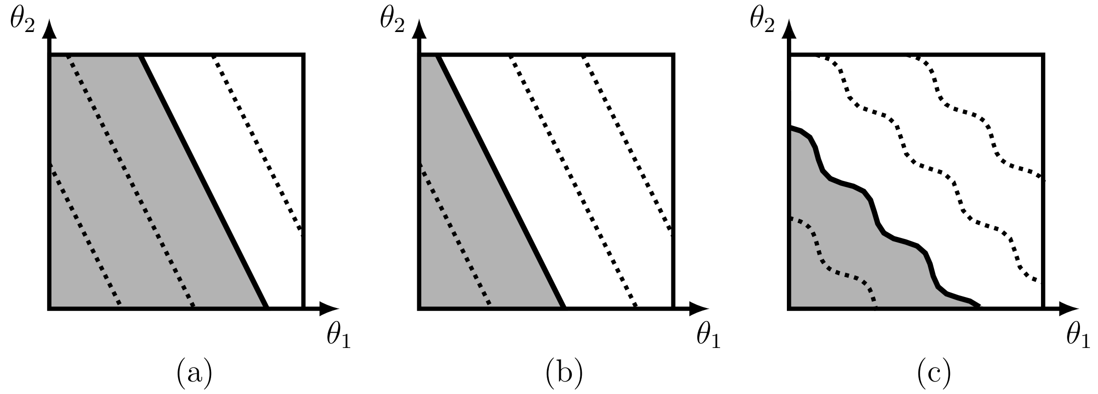

Figure 1: (a) In the efficient allocation, trade occurs below the iso cost curve given by $b = \sigma \cdot \boldsymbol{\theta}$. (b) In the optimal allocation for the single-agent benchmark, trade occurs below the iso-single-agent-virtual-cost curve given by $b = \sigma \cdot \boldsymbol{\theta} + \frac{G(\sigma \cdot \boldsymbol{\theta})}{g(\sigma \cdot \boldsymbol{\theta})}$. (c) In the optimal allocation for our group setting, trade occurs below the iso-group-virtual-cost curve given by $b = \omega \cdot \boldsymbol{\varphi}$.

In a buyer-optimal mechanism for the single-agent benchmark, the good is traded whenever the benefit exceeds its single-agent *virtual* cost,

$$b \geqslant \sigma \cdot \boldsymbol{\theta} + \frac{G(\sigma \cdot \boldsymbol{\theta})}{g(\sigma \cdot \boldsymbol{\theta})}.$$

For any c_0, let the **iso-single-agent-virtual-cost curve for** c_0 be the set of all $\theta \in \Theta$ that have the same virtual cost c_0—that is, satisfying $c_0 = \sigma \cdot \theta + \frac{G(\sigma \cdot \theta)}{g(\sigma \cdot \theta)}$. Then trade occurs below the iso-single-agent-virtual-cost curve for b. This curve is shown for two agents in Panel (b) of Figure 1. Observe that every iso-single-agent-virtual-cost curve is also an iso-cost curve (associated with a lower cost level), so that the former curves are also straight lines with slope $-\frac{\sigma_1}{\sigma_2}$, and the virtual cost also increases as we move in the northeast direction. As is well known, optimal allocations for single-agent settings entail a downward distortion in trade: when $\sigma \cdot \boldsymbol{\theta} < b < \sigma \cdot \boldsymbol{\theta} + \frac{G(\sigma \cdot \boldsymbol{\theta})}{G(\sigma \cdot \boldsymbol{\theta})}$, trade occurs in the efficient allocation but not according to optimal allocations in the single-agent benchmark.

In buyer-optimal mechanisms in our group setting, trade occurs whenever the benefit exceeds the weighted virtual cost of all group members,

$$b \geqslant \omega \cdot \boldsymbol{\varphi} = \sum_i \omega_i \left[\boldsymbol{\theta}_i + \frac{F_i(\boldsymbol{\theta}_i)}{f_i(\boldsymbol{\theta}_i)} \right].$$

For any c_0, let the **iso-group-virtual-cost curve for** c_0 be the set of all type profiles that have the same weighted virtual cost c_0, i.e., those $\theta \in \Theta$ satisfying $c_0 = \sum_i \omega_i \left[\theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)} \right]$. Then, in the optimal allocation rule for our group setting, trade occurs below the iso-group-virtual-cost curve for b. This allocation rule is illustrated for two agents in panel (c) of Figure 1. Notably, because the weights ω

do not depend on the shares σ, iso-group-virtual-cost curves are unrelated to the shares (holding fixed the distribution of $\boldsymbol{\theta}$).

A comparison of the iso-cost curves with the iso-group-virtual-cost curves shows that an optimal allocation rule in our group setting differs from the efficient allocation for *three* reasons. To see this, let us convert the weighted virtual costs to (actual) costs in three steps, and study each conversion:

$$\sum_i \omega_i \left[\boldsymbol{\theta}_i + \frac{F_i(\boldsymbol{\theta}_i)}{f_i(\boldsymbol{\theta}_i)}\right] \rightsquigarrow \sum_i \sigma_i \left[\boldsymbol{\theta}_i + \frac{F_i(\boldsymbol{\theta}_i)}{f_i(\boldsymbol{\theta}_i)}\right] \rightsquigarrow \sum_i \sigma_i \left[\boldsymbol{\theta}_i + \frac{F_i(\sigma \cdot \boldsymbol{\theta})}{f_i(\sigma \cdot \boldsymbol{\theta})}\right] \rightsquigarrow \sum_i \sigma_i \boldsymbol{\theta}_i.$$

The first conversion highlights a *rotational* distortion. Whereas the iso-group-virtual-cost curves are unaffected by the shares σ, the iso-cost curves rotate as the shares change. With two agents, as σ_1 increases, each iso-cost curve rotates clockwise whereas iso-group-virtual-cost curves are unaltered. The second conversion highlights a *curvature* distortion. Unlike iso-cost curves, iso-group-virtual-cost curves might be non-linear because the inverse hazard rate functions F_i/f_i might be non-linear. The third conversion highlights the familiar *downward* distortion. The addition of the inverse hazard rate term elevates the iso-weighted-virtual-cost curves and leads to lower probability of trade. The optimal mechanism for the single-agent benchmark exhibits only the third distortion (with a different inverse hazard rate, G/g) and not the other two.

Two salient features emerge from examining how these three different distortions interact. First, as we will demonstrate, trade in the group setting may be *inefficiently high*. That is, optimal allocations in our main model may prescribe trade even when trading is inefficient. This type of inefficient trade cannot happen in the single-agent benchmark, which exhibits only the downward distortion. Second, an efficiency comparison between buying from a group and buying from a single agent is ambiguous in general. As we will show, a key determinant of the welfare ranking is how large a benefit the good yields for the buyer. Focusing on the natural case in which the per-unit costs $\{\boldsymbol{\theta}_i\}_i$ are identically distributed, we show optimal allocations in our group setting are more efficient than in the single-agent setting if b is large, and are less efficient when b is small. We elaborate more on these two observations below.

For the first observation, suppose each $i \in N$ has $\boldsymbol{\theta}_i \in [0,1]$ following the power distribution $F_i(\theta_i) = \theta_i^\alpha$ for a power $\alpha > 0$. Recall that identically distributed $\{\boldsymbol{\theta}_i\}_i$ lead to $\omega = \left(\frac{1}{N}, \ldots, \frac{1}{N}\right)$ being optimal in the group setting, and that $\varphi_i(\theta_i) = \frac{\alpha+1}{\alpha}\theta_i$ for any $\theta_i \in [0,1]$. Given the latter linear form, iso-weighted-virtual-cost curves are linear, and so there is no curvature distortion in this example. The overall distortion depends on the interaction between rotational and downward distortions, and the rotational distortion might dominate for certain type profiles.

In particular, compare the weighted virtual cost,

$$\omega \cdot \boldsymbol{\varphi} = \sum_i \frac{1}{N} \boldsymbol{\varphi}_i = \frac{1+\alpha}{\alpha N} \sum_i \boldsymbol{\theta}_i,$$

to the (actual) cost,

$$\sum_i \sigma_i \boldsymbol{\theta}_i.$$

Suppose the share vector σ is asymmetric, so that some i has $\sigma_i > \frac{1}{N}$. Then, any large enough α admits a range of b for which

$$\frac{1+\alpha}{\alpha N} < b < \sigma_i.$$

In this case, with positive probability—specifically, when $\boldsymbol{\theta}_i$ is high and $\{\boldsymbol{\theta}_j\}_{j \neq i}$ are low—trading is inefficient but still happens under the optimal allocation rule for the group setting. This example suggests that the irrelevance of land shares, σ, to the optimal mechanism may distort trade in favor of smaller landholders with low productivity.

Second, consider the efficiency of the allocation. Specializing to the case in which $\{\boldsymbol{\theta}_i\}_i$ are identically distributed, we provide an efficiency ranking between the group and single-agent settings for two cases: when the benefit to the buyer is very low, and when it is very high. First, when the benefit is low enough, we show our group setting generates a lower expected surplus than the single-agent benchmark. This surplus ranking holds in an ex-post sense—that is, the buyer's chosen mechanism for the single-agent setting stipulates trade whenever trade is efficient and happens in the group setting—if and only if the shares are sufficiently similar. On the other hand, when the benefit is large enough, then our group setting yields more surplus (in the stronger ex-post sense) than the single-agent one whatever the share vector. In particular, our results imply that the efficiency ranking between the group and single-agent settings will generally depend on the specific parameters of the model. We formally state and prove these results in Appendix B.

To provide some intuition, let us focus on the case in which the land shares are symmetric. In this case, because both weighted virtual costs and single-agent virtual costs are above actual costs, it follows that trading is efficient whenever the optimal allocation in the group or the single-agent setting prescribes it. Therefore, a surplus ranking will follow from showing one regime specifies trade in a bigger region than the other. When the actual cost realizations are extreme—either very high or very low—we can establish this ranking of trade regions. This is because, in these cases, we can rank the single-agent virtual costs against the weighted virtual costs. The case of high costs is simpler: Because the average cost can only be high if *all* sellers have a high cost, the density of the average cost must vanish at the tails, leading to an infinite single-agent virtual cost (whereas the weighted virtual cost is finite). The case of low costs requires a more detailed quantitative calculation, but we show in the appendix that single-agent virtual costs are indeed lower than weighted virtual costs when the average actual cost is low. The ranking follows: when the benefit to the buyer is very low, the single-agent optimal mechanism generates more surplus than the optimal mechanism in the group setting, while the reverse ranking holds when the benefit is high.

7. Discussion

We now consider some variants of our main model and briefly discuss how our analysis can be extended in these directions. Any nontrivial formalism is deferred to Appendix C.

Dominant strategies. The notion of incentive compatibility we have employed so far is Bayesian incentive compatibility (BIC, which we have called IC throughout), which requires only that sellers' reports be best responses in expectation, given their own realized types. However, in our leading application—a group of sellers who collectively own a plot of land—one could envision scenarios without any private information inside the group. That is, the group members might know other members' costs, but the buyer does not.

Motivated by such situations, it is perhaps natural to consider more demanding incentive constraints. Specifically, we consider what happens when the buyer is constrained to offer a mechanism that is dominant-strategy incentive compatible (DIC). Whereas proportional transfers impose no constraints on what allocation rules can be implemented under Bayesian incentive constraints, we show they significantly constrain a buyer restricted to DIC mechanisms; that is, there is no counterpart to part (i) of Lemma 1 saying every ex-post monotone allocation rule is DIC-implementable by some transfer rule.[19] In particular, all deterministic DIC mechanisms take the form of a posted price (augmented by an upfront transfer) with trade occurring if and only if enough sellers approve the trade. Using this observation, we show that no optimal mechanism (i.e., those characterized by Theorem 1) is also DIC. The intuition is similar to that of Proposition 2: optimal mechanisms (putting weight on multiple agents) deliver smooth incentives to a

[19]This fact is reminiscent of other work showing BIC-DIC equivalence fails given financial constraints. Even though unidimensional private-values mechanism design settings with flexible transfers admit a strong form of this equivalence (Gershkov et al., 2013), a DIC constraint severely constrains implementable allocations when paired with ex post budget balance. See (Hagerty and Rogerson, 1987), for example.

single agent, whereas deterministic DIC mechanisms cannot.

Ex-post participation. With a view to respecting individual property rights, we have constrained our buyer to employ a mechanism that is individually rational for each seller—that is, such that every seller can keep his land rather than interacting with the mechanism. As with our other incentive constraints, we formulated IR in the interim sense, having each seller assess his participation decision in expectation over others' types. One may wish to consider a buyer constrained by a stronger form of property rights—namely, that any seller has the option to walk away from the mechanism even after all uncertainty has been resolved. For some examples, such an ex post IR constraint imposes no additional costs on the buyer.[20] For instance, in the equilibrium described in the bidding game of Example 1, no seller ever has an incentive to walk away, even after he learns the other seller's bid if the powers in the sellers' distributions multiply to at most one, $\alpha_1\alpha_2 \leqslant 1$. When does this property hold more generally? And can our analytical approach be applied to understand such ex post constraints?

To better understand the ex post IR constraint more generally, we first answer an implementability question: When can a given allocation rule, together with an expected transfer, be implemented in some IC and ex post IR mechanism? Using this characterization, we provide sufficient conditions on primitives under which some buyer-optimal mechanism is also ex post IR. Applied to Example 1, these conditions say if $\alpha_1 = \alpha_2 \leqslant 1$, some optimal mechanism exists that is also ex post IR, consistent with our calculations for that example.

Pareto-optimal mechanisms. Our paper has focused on mechanisms that maximize the buyer's expected payoff. Although this objective is a natural benchmark, it assumes the buyer has extreme bargaining power relative to the seller group. More generally, one might wonder what mechanisms can arise naturally with different allocations of bargaining rights. Specifically, we study Pareto-optimal mechanisms—that is, IC and IR mechanisms for which there is no alternative IC and IR mechanism that delivers a weakly higher buyer profit, and a weakly higher agent i value for each agent i, with at least one of these $N + 1$ inequalities strict. Our land-acquisition application suggests another reason to care about the entire Pareto frontier. If the buyer is a government, state-owned enterprise, or other large stakeholder in the relevant community, then they may care about the welfare of the current landholders in addition to the purely financial consequences of trade. Understanding the range of optimal mechanisms all such stakeholders might wish to use amounts to understanding all Pareto optima of our space of IC and IR mechanisms.

[20] A similar equivalence arises in related work by Che (2002).

Pareto-optimal mechanisms trade if and only if the buyer's benefit maximizes a weighted average of sellers' *virtual and actual* costs. Although this class of allocation rules is richer than the unique buyer optimum, it enjoys similar tractability and qualitative structure. For instance, Pareto-optimal mechanisms are deterministic and use weights that are fixed and do not depend on reports. The weight that applies to a seller's cost is exactly the Pareto weight of that seller, whereas the weight that applies to the virtual cost is identified endogenously and reflects the agent's influence over the outcomes. We also use our characterization to generalize the main result of Section 5, showing every Pareto-optimal mechanism entails complex pricing.

Our characterization of implementable allocation rules, along with the analytical approach we adopt in developing Theorem 1, proves useful in providing our characterization of Pareto-optimal mechanisms. A standard separation result enables us to represent Pareto optima as maximizers of weighted sums of the $N + 1$ individuals' objectives, and we can adapt our zero-sum game proof to this more general class of objectives.

Pre-market trade. Throughout, we have restricted attention to mechanisms in which agents are paid proportionally to their land shares. We mainly impose this structure as a fairness or institutional requirement that is natural in many applications. As formalized in the appendix, we point out another desirable property of such mechanisms for the case in which sellers' per-unit costs are identically distributed. We show that if the buyer uses these proportional transfer mechanisms, then sellers have no incentives to manipulate the outcome by trading their shares before interacting with the mechanism. We also show, by example, that this property could be violated if the buyer were not restricted to paying agents proportionally to their shares (as in Güth and Hellwig, 1986). Thus, in addition to being realistic in many settings, our assumption of collective transfers yields a desirable robustness property for buying mechanisms.

The result that the sellers have no incentives to trade shares is based on two observations. First, the optimal mechanism is independent of the shares. Second, in a mechanism that is independent of the shares, the sellers' incentives to trade shares disappear if agents are paid the same price per share. When discriminatory pricing is allowed, the buyer optimally treats sellers with different shares differently, opening the door to gaming by trading shares.

Beyond veto bargaining. An important feature of our environment is that any agent can unilaterally veto the mechanism. This feature, captured by the requirement that the mechanism be IR for all of the agents, is natural in settings with strong property rights. However, a more permissive bargaining arrangement may be more appropriate for modeling some contexts—for example, when eminent domain enables a government to forcibly acquire land from some individuals for public projects. As in some redevelopment projects, we could require that rather than unanimity, the terms of trade need to be approved by at least n agents for some given $n < N$. This flexibility raises new modeling questions concerning how exactly one determines whether a mechanism has sufficient approval.

In one approach, we could require this approval by n sellers determined ex ante—that is, independent of their type realizations. This formulation reduces nearly immediately to the analysis in our main model. Indeed, one need only replace the IR constraint (which we imposed for all N agents in our model) with a weaker assumption that at least n agents' IR constraints are satisfied. Because the buyer has no reason to condition on the types of agents facing no IR constraint, her problem reduces to an n-agent specification of our main model. The optimal mechanism allocates the good if and only if the benefit to the buyer exceeds a weighted sum of the chosen n agents' virtual costs. The buyer would then choose to tailor the mechanism to the n agents she finds most favorable to interact with ex ante—for instance (given Theorem 1), the n agents with the lowest virtual cost distributions if these distributions are first-order stochastically ranked.

References

L. M. Ausubel, P. Cramton, M. Pycia, M. Rostek, and M. Weretka. Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies*, 81(4):1366–1400, 2014. 5

Y. Azrieli and S. Kim. Pareto efficiency and weighted majority rules. *International Economic Review*, 55(4):1067–1088, 2014. 6

M. Babaioff, N. Immorlica, B. Lucier, and S. M. Weinberg. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)*, 67(4):1–40, 2020. 19

A. V. Banerjee. Land reforms: prospects and strategies. *Available at SSRN 183711*, 1999. 4, 18

R. E. Barlow, A. W. Marshall, and F. Proschan. Properties of probability distributions with monotone hazard rate. *The Annals of Mathematical Statistics*, pages 375–389, 1963. 21

T. C. Bergstrom. Cournot equilibrium in factor markets. 1978. 6

R. A. Berry, W. R. Cline, et al. *Agrarian structure and productivity in developing countries: a study prepared for the International Labour Office within the* *framework of the World Employment Programme.* Johns Hopkins Univ. Press, 1979. 18

T. Börgers and P. Norman. A note on budget balance under interim participation constraints: the case of independent types. *Economic Theory*, 39(3):477–489, 2009. 11

Y. Cai, C. Daskalakis, and S. M. Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 618–627. IEEE, 2013. 6

R. Chand, P. L. Prasanna, and A. Singh. Farm size and productivity: Understanding the strengths of smallholders and improving their livelihoods. *Economic and Political Weekly*, pages 5–11, 2011. 2

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010. 19

S. Chawla, D. Malec, and B. Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015. 19

Y.-K. Che. The economics of collective negotiation in pretrial bargaining. *International Economic Review*, pages 549–575, 2002. 7, 26

Y.-K. Che and J. Kim. Robustly collusion-proof implementation. *Econometrica*, 74(4):1063–1107, 2006. 11

P. Cramton, R. Gibbons, and P. Klemperer. Dissolving a partnership efficiently. *Econometrica: Journal of the Econometric Society*, pages 615–632, 1987. 6

C. d'Aspremont and L.-A. Gérard-Varet. Incentives and incomplete information. *Journal of Public economics*, 11(1):25–45, 1979. 6

C. d'Aspremont, J. Crémer, and L.-A. Gérard-Varet. Balanced bayesian mechanisms. *Journal of Economic Theory*, 115(2):385–396, 2004. 11

M. Ekmekci, N. Kos, and R. Vohra. Just enough or all: Selling a firm. *American Economic Journal: Microeconomics*, 8(3):223–256, 2016. 6

M. Friedman. *A program for monetary stability.* Number 3. Ravenio Books, 1960. 5

S. Gandhi, V. Tandel, A. Tabarrok, and S. Ravi. Too slow for the urban march: Litigations and the real estate market in mumbai, india. *Journal of Urban Economics*, 123:103330, 2021. 2

A. Gershkov, J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi. On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1):197–220, 2013. 25

A. Gershkov, B. Moldovanu, and X. Shi. Optimal voting rules. *The Review of Economic Studies*, 84(2):688–717, 2017. 6

P. Gopalan, N. Nisan, and T. Roughgarden. Public projects, boolean functions, and the borders of border's theorem. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):1–21, 2018. 11

S. J. Grossman and O. D. Hart. Takeover bids, the free-rider problem, and the theory of the corporation. *The Bell Journal of Economics*, pages 42–64, 1980. 7

W. Güth and M. Hellwig. The private supply of a public good. *Journal of Economics*, 46(1):121–159, 1986. 6, 27, 22

K. M. Hagerty and W. P. Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, 42(1):94–107, 1987. 25

S. Hart and N. Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017. 19

M. A. Heller. The tragedy of the anticommons: property in the transition from marx to markets. *Harvard law review*, pages 621–688, 1998. 2

R. Kirkegaard. A mechanism design approach to ranking asymmetric auctions. *Econometrica*, 80(5):2349–2364, 2012. 16

V. Krishna and J. Morgan. Majority rule and utilitarian welfare. *American Economic Journal: Microeconomics*, 7(4):339–75, 2015. 6

T. Kuran. Why the middle east is economically underdeveloped: historical mechanisms of institutional stagnation. *Journal of economic perspectives*, 18(3):71–90, 2004. 2

G. J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents. *The Review of Economic Studies*, 57(3):351–367, 1990. 6

L. Makowski and C. Mezzetti. Bayesian and weakly robust first best mechanisms: characterizations. *Journal of Economic Theory*, 64(2):500–519, 1994. 11

P. F. Malvey and C. M. Archibald. Uniform-price auctions: Update of the treasury experience. *US Treasury*, 1998. 5

A. Manjunatha, A. R. Anik, S. Speelman, and E. Nuppenau. Impact of land fragmentation, farm size, land ownership and crop diversity on profit and efficiency of irrigated farms in india. *Land use policy*, 31:397–405, 2013. 2

E. Maskin and J. Riley. Asymmetric auctions. *The review of economic studies*, 67 (3):413–438, 2000. 16

R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6 (1):58–73, 1981. 3, 10, 18, 32, 19

R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983. 10, 32, 19

S. Oliveros and M. Iaryczower. Collective hold-up. *Theoretical Economics*, 2022. 7

M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT press, 1994. 33

D. W. Rae. Decision-rules and individual values in constitutional choice. *The American Political Science Review*, 63(1):40–56, 1969. 6

J. Riley and R. Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289, 1983. 18

R. Rob. Pollution claim settlements under private information. *Journal of Economic Theory*, 47(2):307–333, 1989. 6

P. W. Schmitz and T. Tröger. The (sub-) optimality of the majority rule. *Games and Economic Behavior*, 74(2):651–665, 2012. 6

M. Shaked and J. G. Shanthikumar. *Stochastic orders*. Springer Science & Business Media, 2007. 39, 40

A. Sood. Land market frictions in developing countries: Evidence from manufacturing firms in india. Technical report, 2020. 2

Appendix

A. Proofs for main results

A.1. Proofs for Section 3

We now reproduce the statement of Lemma 1.

Lemma: *Let x be some allocation rule.*

(i) Mechanism (x, m) is IC and IR for some transfer rule m if and only if x is interim monotone.

(ii) If some transfer rule m exists such that mechanism (x, m) is IC and IR, then a maximally profitable such mechanism exists, with resulting profit

$$\min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right].$$

Proof of Lemma 1. For each $i \in N$, let $X_i := X_i^x$, and define $M_i^* : \Theta_i \to \mathbb{R}$ by

$$M_i^*(\theta_i) := X_i(\theta_i)\theta_i + \int_{\theta_i}^{\bar{\theta}_i} X_i(\tilde{\theta}_i) \, \mathrm{d}\tilde{\theta}_i.$$

Given a transfer rule m, standard arguments (Myerson, 1981; Myerson and Satterthwaite, 1983) show that (x, m) is IC if and only if each $i \in N$ has X_i weakly decreasing and $M_i^m = M_i^* + \underline{U}_i$ for some constant $\underline{U}_i \in \mathbb{R}$; that such a mechanism is IR if and only if $\underline{U}_i \geqslant 0$ for each $i \in N$; and that $\mathbb{E}\left[M_i^*(\boldsymbol{\theta}_i)\right] = \mathbb{E}\left[X_i(\boldsymbol{\theta}_i)\boldsymbol{\varphi}_i\right]$. Given iterated expectations, the latter equation simplifies to $\mathbb{E}\left[M_i^*(\boldsymbol{\theta}_i)\right] = \mathbb{E}\left[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i\right]$.

Using the above observations, let us prove the two parts of the lemma in turn.

Toward part (i), note the first paragraph says interim monotonicity is necessary for x to be IC implementable; and for sufficiency it suffices to show some transfer rule m^0 exists such that $M_i^{m^0} - M_i^*$ is constant for each $i \in N$ (since raising such a transfer rule by a large enough constant will ensure IR). The transfer rule m^0 given by $m^0(\theta) := \sum_{i \in N} M_i^*(\theta_i)$ has this property, and so part (i) follows.

Now, toward part (ii), suppose x is indeed implementable; say transfer rule m is such that (x, m) is IC and IR. Then each $i \in N$ admits $\underline{U}_i \geqslant 0$ such that $M_i^m = M_i^* + \underline{U}_i$. Hence, for any $i \in N$, we can write the expected transfer as

$$\mathbb{E}\left[m(\boldsymbol{\theta})\right] = \mathbb{E}\left[M_i^M(\boldsymbol{\theta}_i)\right] = \mathbb{E}\left[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i\right] + \underline{U}_i,$$

so that the buyer's expected value can be written as

$$\mathbb{E}\left[bx(\boldsymbol{\theta}) - m(\boldsymbol{\theta})\right] = \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] - \underline{U}_i.$$

Reducing the transfer rule by a constant will reduce each of $\{\underline{U}_i\}_i$ by the same constant, and so raise the buyer's expected value. The buyer therefore optimally sets $\min_{i \in N} \underline{U}_i = 0$. But in this case, we have $\mathbb{E}\left[bx(\boldsymbol{\theta}) - m(\boldsymbol{\theta})\right] \leqslant \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right]$ for every $i \in N$, with equality for some i. Said differently, we then have

$$\mathbb{E}\left[bx(\boldsymbol{\theta}) - m(\boldsymbol{\theta})\right] = \min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right],$$

delivering part (ii) □

The following notation will be convenient to us in making formal arguments.

Notation 1: *Let $\tilde{\mathcal{X}}$ denote the set of all allocation rules $\mathcal{X}$, modulo the F-almost everywhere equivalence relation, a subset of $L^\infty(\Theta, F)$. Each element of $\mathcal{X}$ corresponds to one of $\tilde{\mathcal{X}}$ in the obvious way.*

Consider now the relaxed buyer problem,

$$\max_{x \in \tilde{\mathcal{X}}} \left\{ \min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] \right\}. \tag{RBP}$$

which is our buyer's problem without the interim-monotonicity constraint (and cast in $\tilde{\mathcal{X}}$). The following lemma characterizes solutions of this relaxed program.

Lemma 2: *A unique solution exists to program (RBP). This solution is given by the ω-allocation rule, where $\omega \in \Delta N$ is any weight vector satisfying the two equivalent conditions (i) and (ii) in the statement of Theorem 1.*

Proof. Consider a two-player zero-sum game where the maximizer (Max) chooses $x \in \tilde{\mathcal{X}}$ and the minimizer (Min) chooses $\omega \in \Delta N$. The objective (that is, the payoff to Max) is

$$\mathcal{G}(x, \omega) := \mathbb{E}[x(\boldsymbol{\theta})(b - \omega \cdot \boldsymbol{\varphi})].$$

We will first argue that a Nash equilibrium exists for this zero-sum game; and that the Nash equilibria are exactly the pairs $(x^*, \omega^*) \in \tilde{\mathcal{X}} \times \Delta N$ for which x^* solves (RBP) and ω^* satisfies condition (i). Then we will argue that x_ω is Max's unique best response to any Min strategy ω; that Max has a unique Nash equilibrium strategy; and that condition (ii) is equivalent to being a Nash equilibrium strategy for Min. Establishing these facts will establish the lemma.

First, because $\tilde{\mathcal{X}}$ is weak*-compact (by Banach-Alaoglu) and convex, the space ΔN obviously is as well, and the objective is weak*-continuous in the strategy profile, it follows from Sion's minimax theorem that

$$\max_{x \in \tilde{\mathcal{X}}} \min_{\omega \in \Delta N} \mathcal{G}(x, \omega) = \min_{\omega \in \Delta N} \max_{x \in \tilde{\mathcal{X}}} \mathcal{G}(x, \omega),$$

where all extrema in the equation are attained by Berge's theorem. Then, because the auxiliary game is strictly competitive, Proposition 22.2 from Osborne and Rubinstein (1994) tells us some Nash equilibrium exists, and that the Nash equilibria are exactly the pairs $(x^*, \omega^*) \in \tilde{\mathcal{X}} \times \Delta N$ for which

$$\begin{aligned} x^* &\in \operatorname{argmax}_{x \in \tilde{\mathcal{X}}} \min_{\omega \in \Delta N} \mathcal{G}(x, \omega) \text{ and} \\ \omega^* &\in \operatorname{argmin}_{\omega \in \Delta N} \max_{x \in \tilde{\mathcal{X}}} \mathcal{G}(x, \omega). \end{aligned}$$

(In particular, the set of equilibria forms a product set.) Observe, though, that $\min_{\omega \in \Delta N} \mathcal{G}(x, \omega) = \min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right]$ for each $x \in \tilde{\mathcal{X}}$ because $\mathcal{G}(x, \cdot)$ affine. Hence, x^* maximizes this quantity if and only if x^* solves (RBP). Moreover, $\max_{x \in \tilde{\mathcal{X}}} \mathcal{G}(x, \omega) = \mathbb{E}\left[\max_{\mathtt{x} \in [0,1]} (b - \omega \cdot \boldsymbol{\varphi})\mathtt{x}\right] = \mathbb{E}\left[(b - \omega \cdot \boldsymbol{\varphi})_+\right]$ for each $\omega \in \Delta N$, so that minimizing these expressions is equivalent—that is, the minimax strategies are exactly those satisfying condition (i). So we have established that some Nash equilibrium exists; and that the Nash equilibria are exactly the pairs $(x^*, \omega^*) \in \tilde{\mathcal{X}} \times \Delta N$ for which x^* solves (RBP) and ω^* satisfies condition (i).

It remains to show that x_ω is Max's unique best response to any Min strategy ω; that Max has a unique Nash equilibrium strategy; and that condition (ii) is equivalent to being a Nash equilibrium strategy for Min. Toward the first assertion, consider any $\omega \in \Delta N$. Because $\{\boldsymbol{\theta}_i\}_{i\in N}$ are atomless and independent and $\{\varphi_i\}_{i\in N}$ are all strictly increasing, it follows that $\mathbb{P}\{\omega \cdot \boldsymbol{\varphi} = b\} = 0$, so that the $\tilde{\mathcal{X}}$ element with representative x_ω is the unique $x \in \tilde{\mathcal{X}}$ such that

$$\mathbb{P}\left\{x(\boldsymbol{\theta}) \in \operatorname{argmax}_{\mathrm{x}\in[0,1]}\left[(b - \omega \cdot \boldsymbol{\varphi})\mathrm{x}\right]\right\} = 1.$$

Thus, the ω-allocation rule x_ω is Min's unique best response (in $\tilde{\mathcal{X}}$) to ω. From the product structure of the set of Nash equilibria, then, it follows that Max has a unique Nash equilibrium strategy x^*, which is then the unique solution to (RBP).

All that remains now is to show that condition (ii) is equivalent to being a Nash equilibrium strategy for Min. But because x_ω is the unique Max best response to $\omega \in \Delta N$, we know ω is a Nash equilibrium strategy if and only $\omega \in \operatorname{argmax}_{\tilde{\omega}\in\Delta N} \mathcal{G}(x_\omega, \tilde{\omega})$ or, equivalently (since $\mathcal{G}(x_\omega, \cdot)$ is affine) every $i \in \operatorname{supp}(\omega)$ belongs to $\operatorname{argmax}_{i\in N} \mathbb{E}\left[(b - \boldsymbol{\theta}_i)\, \mathbb{1}_{\omega\cdot\boldsymbol{\theta}\leqslant b}\right]$. Finally $b > \underline{\theta}_i$ for every $i \in N$, the event $\{\omega\cdot\boldsymbol{\theta} \leqslant b\}$ has strictly positive probability, so that the latter condition is equivalent to condition (i). The lemma follows. □

We now reproduce the statement of Theorem 1.

Theorem (Optimal allocation)**:**
A weighted allocation rule is essentially uniquely optimal. The unique optimal-weight vector ω is characterized by either of the following two equivalent conditions:

(i) $\omega \in \operatorname{argmin}_{\tilde{\omega}\in\Delta N} \mathbb{E}[(b - \tilde{\omega} \cdot \boldsymbol{\varphi})_+]$.

(ii) $\operatorname{supp}(\omega) \subseteq \operatorname{argmax}_{i\in N} \mathbb{E}\left[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b\right]$.

Moreover, if $b < \bar{\theta}_j$ for at least two $j \in N$, then every $i \in N$ has $\omega_i < 1$.

Proof of Theorem 1. First, by Lemma 1, an allocation rule is optimal if and only if it solves program (BP), so we focus on solutions to this program.

Now, Lemma 2 tells us that conditions (i) and (ii) in the theorem's statement are equivalent, that some $\omega \in \Delta N$ exists that satisfies those conditions, and that (the almost-sure equivalence class of) x_ω is uniquely optimal in (RBP). Because x_ω is interim monotone and solves a relaxation of (BP), it follows directly that x_ω solves (BP), and that every other solution x to (BP) has $x(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely.

Next, we establish $i \in N$ has $\omega_i < 1$ if some $j \in N\backslash\{i\}$ has $b < \bar{\theta}_j$. To see this fact, note that if $i \in N$ had $\omega_i = 1$, then $j \in N\backslash\{i\}$ would have

$$\mathbb{E}\left[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b\right] = \mathbb{E}\left[\boldsymbol{\varphi}_i \mid \boldsymbol{\varphi}_i \leqslant b\right] \leqslant b < \bar{\theta}_j = \mathbb{E}\left[\boldsymbol{\varphi}_j\right] = \mathbb{E}\left[\boldsymbol{\varphi}_j \mid \omega \cdot \boldsymbol{\varphi} \leqslant b\right],$$

in contradiction to condition (ii).

Finally, we turn to uniqueness of ω. Suppose $\tilde{\omega} \in \Delta N$ is such that $x_{\tilde{\omega}}$ is optimal, and so $x_{\tilde{\omega}}(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely; our aim is to show $\tilde{\omega} = \omega$. Toward establishing this equality, define $G := \prod_{i\in N} \left(b - \varphi_i(\bar{\theta}_i),\ b - \underline{\theta}_i\right)$, the interior of the support of $b\mathbb{1}_N - \boldsymbol{\varphi}$. Now, define the linear map $L : \mathbb{R}^N \to \mathbb{R}^2$ by letting $L(z) := (\omega \cdot z,\ \tilde{\omega} \cdot z)$ for each $z \in \mathbb{R}^N$. Let us now observe some properties of G and L. First, that $\omega \cdot \varphi(\underline{\theta}) < b < \omega \cdot \varphi(\bar{\theta})$ and $\tilde{\omega} \cdot \varphi(\underline{\theta}) < b < \tilde{\omega} \cdot \varphi(\bar{\theta})$ implies $L(G)$ is not a subset of $\mathbb{R}_+ \times \mathbb{R}$, of $\mathbb{R}_- \times \mathbb{R}$, of $\mathbb{R} \times \mathbb{R}_+$, or of $\mathbb{R} \times \mathbb{R}_-$. Second, that $\mathbb{P}\left\{x_{\tilde{\omega}}(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})\right\} = 1$ implies $L(G)$ is a subset of $\mathbb{R}^2_+ \cup \mathbb{R}^2_-$. Third, because L is linear and G is convex, the set $L(G)$ is convex. Combining these three observations tells us that $L(G)$ is contained in a single line through the origin. Because G is open and L is linear, then, $L(\mathbb{R}^N)$ is contained in the same line. Said differently, the rank of the linear map L is 1, so that vectors $\omega, \tilde{\omega} \in \mathbb{R}^N_+$ are proportional. Because $||\omega||_1 = 1 = ||\tilde{\omega}||_1$, it follows that $\omega = \tilde{\omega}$. □

Proof for Example 1. In what follows, we proceed in three steps. First, we show ω^* is the optimal weight vector, thus characterizing optimal allocation rules. Second, we name a specific (ex-post) transfer rule, and show that this transfer rule paired with our optimal allocation rule constitutes an optimal mechanism. Third, we show that the strategy profile we have named for the bidding game is an equilibrium that induces this optimal mechanisms.

We first prove that ω^* is the optimal vector of weights. To this end, let

$$\nu_i := 1 + \frac{1}{\alpha_i} = \frac{\alpha_i + 1}{\alpha_i}$$

for each agent i, and notice that $\varphi_i(\theta_i) = \nu_i\theta_i$ for every $\theta_i \in \Theta_i = [0, 1]$. For any $\omega \in \Delta N$ with $\min\{\omega_1\nu_1,\ \omega_2\nu_2\} \geqslant b$ (an interval of ω including ω^*), observe that

$$\mathbb{E}[(b - \omega\cdot\boldsymbol{\varphi})_+] = \int_0^{\frac{b}{\omega_1\nu_1}} \int_0^{\frac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}} \left[b - \omega_1\nu_1\theta_1 - (1-\omega_1)\nu_2\theta_2\right] \alpha_1\theta_1^{\alpha_1-1}\alpha_2\theta_2^{\alpha_2-1}\, \mathrm{d}\theta_2\, \mathrm{d}\theta_1.$$

Therefore,

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\omega_1}\mathbb{E}\left\{[b - (\omega_1, 1-\omega_1)\cdot\boldsymbol{\varphi}]_+\right\} &= \int_0^{\frac{b}{\omega_1\nu_1}} \alpha_1\alpha_2 \int_0^{\frac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}} (\nu_2\theta_2 - \nu_1\theta_1)\theta_1^{\alpha_1-1}\theta_2^{\alpha_2-1}\, \mathrm{d}\theta_2\, \mathrm{d}\theta_1 \\
&= \int_0^{\frac{b}{\omega_1\nu_1}} \alpha_1\alpha_2 \left\{ \nu_2\theta_1^{\alpha_1-1}\frac{\left[\frac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}\right]^{\alpha_2+1}}{1+\alpha_2} - \nu_1\theta_1^{\alpha_1}\frac{\left[\frac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}\right]^{\alpha_2}}{\alpha_2} \right\} \mathrm{d}\theta_1 \\
&= \int_0^{\frac{b}{\omega_1\nu_1}} \theta_1^{\alpha_1-1}\left[\tfrac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}\right]^{\alpha_2}\left[\alpha_1\tfrac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2} - \alpha_2\tfrac{\omega_1^*\nu_1}{1-\omega_1^*}\theta_1\right] \mathrm{d}\theta_1.
\end{aligned}$$

Note now that if $\omega_1 = \omega_1^*$, the integrand is then equal to

$$\frac{\mathrm{d}}{\mathrm{d}\theta_1}\left\{\theta_1^{\alpha_1}\left[\tfrac{b-\omega_1\nu_1\theta_1}{(1-\omega_1)\nu_2}\right]^{\alpha_2+1}\right\},$$

so that

$$\frac{\mathrm{d}}{\mathrm{d}\omega_1}\mathbb{E}\left\{[b - (\omega_1, 1-\omega_1)\cdot\boldsymbol{\varphi}]_+\right\} = \left(\tfrac{b}{\omega_1\nu_1}\right)^{\alpha_1} 0^{\alpha_2+1} - 0^{\alpha_1}\left[\tfrac{b}{(1-\omega_1)\nu_2}\right]^{\alpha_2+1} = 0.$$

Hence, ω^* solves the convex program from Theorem 1(i), meaning it is optimal.

Now, consider the mechanism (x, m) given by $x(\theta) := x_{\omega^*}(\theta)$ and

$$m(\theta) := x(\theta)[\kappa b + \beta_1\theta_1 + \beta_2\theta_2],$$

where

$$\begin{aligned} \kappa &= \tfrac{\alpha_1+\alpha_2}{(\alpha_1+1)(\alpha_2+1)}(1 - \alpha_1\alpha_2) \\ \beta_i &= \tfrac{\alpha_i+1}{\alpha_{-i}+1}\alpha_{-i} \text{ for } i \in N. \end{aligned}$$

Let us argue that (x, m) is an optimal mechanism. For each $i \in N$, define the interim allocation rule $X_i := X_i^x$, the interim transfer rule $M_i := M_i^m$, and the interim transfer rule M_i^* as defined in the proof of Lemma 1. If $M_i = M_i^*$ for both $i \in N$, then as explained in Lemma 1 proof, the mechanism (x, m) is IC and has binding IR for both agents, and is therefore best for the buyer among all IC and IR mechanisms with allocation rule x; because x is optimal, it will then follow that (x, m) is optimal. So we now turn to showing $M_i = M_i^*$ for both $i \in N$. To that end, let

$$\gamma_i := \omega_i^*\nu_i = \frac{\alpha_i + 1}{\alpha_1 + \alpha_2} > 0,$$

and note that $b \leqslant \gamma_i$ by hypothesis. Therefore, X_i is zero on $\left(\frac{b}{\gamma_i}, 1\right]$, so that any $\theta_i \in [0, 1]$ has

$$X_i(\theta_i) = \mathbb{P}\left[\gamma_{-i}\boldsymbol{\theta}_{-i} \leqslant b - \gamma_i\theta_i\right] = \left(\tfrac{b-\gamma_i\theta_i}{\gamma_{-i}}\right)_+^{\alpha_{-i}}$$

and

$$\begin{aligned} \int_{\theta_i}^1 X_i &= \mathbb{1}_{\theta_i \leqslant \frac{b}{\gamma_i}} \int_{\theta_i}^{\frac{b}{\gamma_i}} \left(\tfrac{b-\gamma_i\tilde{\theta}_i}{\gamma_{-i}}\right)^{\alpha_{-i}} \, \mathrm{d}\tilde{\theta}_i \\ &= \mathbb{1}_{\theta_i \leqslant \frac{b}{\gamma_i}} \int_{\theta_i}^{\frac{b}{\gamma_i}} \left(\tfrac{b-\gamma_i\tilde{\theta}_i}{\gamma_{-i}}\right)^{\alpha_{-i}} \, \mathrm{d}\tilde{\theta}_i \\ &= \tfrac{\gamma_{-i}}{\gamma_i}\mathbb{1}_{\theta_i \leqslant \frac{b}{\gamma_i}} \int_0^{\frac{b-\gamma_i\theta_i}{\gamma_{-i}}} y^{\alpha_{-i}} \, \mathrm{d}y \\ &= \tfrac{\alpha_{-i}+1}{\alpha_i+1}\mathbb{1}_{\theta_i \leqslant \frac{b}{\gamma_i}} \cdot \tfrac{1}{\alpha_{-i}+1}\left(\tfrac{b-\gamma_i\theta_i}{\gamma_{-i}}\right)^{\alpha_{-i}+1} \\ &= \tfrac{1}{\alpha_i+1}\left(\tfrac{b-\gamma_i\theta_i}{\gamma_{-i}}\right) X_i(\theta_i) \\ &= \left[\tfrac{1}{\alpha_i+1}\tfrac{1}{\gamma_{-i}}b - \tfrac{1}{\alpha_i+1}\tfrac{\gamma_i}{\gamma_{-i}}\theta_i\right] X_i(\theta_i) \\ &= \left[\tfrac{\alpha_1+\alpha_2}{(\alpha_1+1)(\alpha_2+1)}b - \tfrac{1}{\alpha_{-i}+1}\theta_i\right] X_i(\theta_i). \end{aligned}$$

Hence,

$$M_i^*(\theta_i) \;=\; X_i(\theta_i)\theta_i + \int_{\theta_i}^{\bar{\theta}_i} X_i = \left[\tfrac{\alpha_1+\alpha_2}{(\alpha_1+1)(\alpha_2+1)}b + \tfrac{\alpha_{-i}}{\alpha_{-i}+1}\theta_i\right] X_i(\theta_i).$$

Next, note that each $y \in [0,1]$ has

$$\mathbb{E}\left[\boldsymbol{\theta}_{-i}\mathbb{1}_{\boldsymbol{\theta}_{-i}\leqslant y}\right] = \int_0^y \theta_{-i}\left(\alpha_{-i}\theta_{-i}^{\alpha_{-i}-1}\right)\, \mathrm{d}\theta_{-i} = \tfrac{\alpha_{-i}}{\alpha_{-i}+1}y^{\alpha_{-i}+1}.$$

Therefore, each $i \in N$ and $\theta_i \in [0,1]$ has

$$\begin{aligned}
\mathbb{E}\left[\boldsymbol{\theta}_{-i}x(\theta_i,\boldsymbol{\theta}_{-i})\right] &= \mathbb{E}\left[\boldsymbol{\theta}_{-i}\mathbb{1}_{\boldsymbol{\theta}_{-i}\leqslant\frac{b-\gamma_i\theta_i}{\gamma_{-i}}}\right] \\
&= \mathbb{1}_{\theta_i\leqslant\frac{b}{\gamma_i}}\cdot\tfrac{\alpha_{-i}}{\alpha_{-i}+1}\left(\tfrac{b-\gamma_i\theta_i}{\gamma_{-i}}\right)^{\alpha_{-i}+1} \\
&= \tfrac{\alpha_{-i}}{\alpha_{-i}+1}\left(\tfrac{b-\gamma_i\theta_i}{\gamma_{-i}}\right)X_i(\theta_i) \\
&= \alpha_{-i}\tfrac{\alpha_i+1}{\alpha_{-i}+1}\left[\tfrac{\alpha_1+\alpha_2}{(\alpha_1+1)(\alpha_2+1)}b - \tfrac{1}{\alpha_{-i}+1}\theta_i\right]X_i(\theta_i) \\
&= \tfrac{\alpha_{-i}}{(\alpha_{-i}+1)^2}\left[(\alpha_1+\alpha_2)b-(\alpha_i+1)\theta_i\right]X_i(\theta_i).
\end{aligned}$$

It follows that

$$\begin{aligned}
M(\theta_i) &= \mathbb{E}\left\{\left[\kappa b+\beta_i\theta_i+\beta_{-i}\boldsymbol{\theta}_{-i}\right]x(\theta_i,\boldsymbol{\theta}_{-i})\right\} \\
&= \left\{\kappa b+\beta_i\theta_i+\beta_{-i}\tfrac{\alpha_{-i}}{(\alpha_{-i}+1)^2}\left[(\alpha_1+\alpha_2)b-(\alpha_i+1)\theta_i\right]\right\}X_i(\theta_i) \\
&= \left\{\left[\kappa+\beta_{-i}\tfrac{\alpha_{-i}}{(\alpha_{-i}+1)^2}(\alpha_1+\alpha_2)\right]b+\left[\beta_i-\beta_{-i}\tfrac{\alpha_{-i}}{(\alpha_{-i}+1)^2}(\alpha_i+1)\right]\theta_i\right\}X_i(\theta_i) \\
&= \left\{\left[\kappa+\tfrac{\alpha_1\alpha_2}{(1+\alpha_1)(1+\alpha_2)}(\alpha_1+\alpha_2)\right]b+\left[\beta_i-\tfrac{\alpha_1\alpha_2}{(1+\alpha_1)(1+\alpha_2)}(\alpha_i+1)\right]\theta_i\right\}X_i(\theta_i) \\
&= \left[\tfrac{\alpha_1+\alpha_2}{(1+\alpha_1)(1+\alpha_2)}b+\left(\beta_i-\tfrac{\alpha_1\alpha_2}{1+\alpha_{-i}}\right)\theta_i\right]X_i(\theta_i) \\
&= \left[\tfrac{\alpha_1+\alpha_2}{(1+\alpha_1)(1+\alpha_2)}b+\tfrac{\alpha_{-i}}{1+\alpha_{-i}}\theta_i\right]X_i(\theta_i) \\
&= M_i^*(\theta_i).
\end{aligned}$$

Hence, the given mechanism (x,m) is optimal.

Finally, let us turn to the bidding game in which each agent i can submit any bid $s_i \geqslant 0$; trade occurs if and only if

$$\tau_1 s_1 + \tau_2 s_2 \leqslant b,$$

where

$$\tau_i := \frac{\alpha_{-i}+1}{(\alpha_1+\alpha_2)\alpha_{-i}} \text{ for } i \in N;$$

and the price (paid if and only if trade occurs) is $p = \kappa b + s_1 + s_2$. Let us consider the strategy profile in which each type θ_i of each agent i bids $\beta_i\theta_i$. We will argue that this strategy profile constitutes a Bayes Nash equilibrium and that it generates allocation rule x and transfer rule m; optimality will then follow from optimality of the mechanism (x,m). First, any type profile $\theta \in \Theta$ has

$$\tau_i\beta_i\theta_i = \omega_i^*\nu_i\theta_i = \omega_i^*\varphi_i(\theta_i)\ \forall i \in N \implies \tau_1\beta_1\theta_1+\tau_2\beta_2\theta_2 = \omega^*\cdot\varphi(\theta).$$

Because each agent i bids $s_i = \beta_i\theta_i$, it follows that trade occurs if and only if $\omega^* \cdot \varphi(\theta) \leqslant b$—that is, the induced allocation rule is exactly x. Second, if trade happens at type profile θ, the price paid under this strategy profile is

$$p = \kappa b + s_1 + s_2 = \kappa b + \beta_1\theta_1 + \beta_2\theta_2.$$

Thus, the induced transfer rule is exactly m. All that remains then is to check that the described bidding rule is an equilibrium. To that end, consider any type $\theta_i \in [0,1]$ of any agent i; we want to show $\beta_i\theta_i$ yields a weakly higher expected payoff for this type than any other $s_i \geqslant 0$. Because the mechanism (x,m) is IC, we know that $\theta_i \in \operatorname{argmax}_{\tilde{\theta}_i\in[0,1]} \mathbb{E}\left[m(\tilde{\theta}_i, \boldsymbol{\theta}_{-i}) - \theta_i x(\tilde{\theta}_i, \boldsymbol{\theta}_{-i})\right]$. But then, because the bidding game and strategy profile induce x and m, it follows that θ_i has no profitable deviation in $\left\{\beta_i\tilde{\theta}_i\right\}_{\tilde{\theta}_i\in[0,1]} = [0,\beta_i]$. Meanwhile, every bid $s_i \geqslant \beta_i$ has

$$\tau_i s_i \geqslant \tau_i\beta_i = \gamma_i \geqslant b,$$

and so (because agent $-i$ has a strictly positive bid almost surely) leads to a zero probability of trade. Thus, all bids $s_i \geqslant \beta_i$ are payoff-equivalent for i, and so do not constitute profitable deviations because bid β_i does not. Hence, the given strategy profile is an equilibrium, as required. □

A.2. Proofs for Section 4

The following lemma provides a sufficient condition to be able to weakly rank agents' weights in the optimal mechanism, and further provides a quantitative sufficient conditions for one agent's weight to be substantially higher than another's.

Lemma 3: *Suppose constants $\alpha \in (0,1]$ and $\beta \geqslant (1-\alpha)\frac{\bar{\theta}_i+\alpha\bar{\theta}_j}{1+\alpha}$ are such that*

$$\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \alpha\boldsymbol{\varphi}_j + \beta.$$

Then, the optimal weight vector ω satisfies $\alpha\omega_i \geqslant \omega_j$.

Proof. Suppose $\boldsymbol{\varphi}_i, \boldsymbol{\varphi}_j, \alpha, \beta$ satisfy the given hypotheses, and let $\omega \in \Delta N$ have $\alpha\omega_i < \omega_j$ (which in particular implies $\omega_j > 0$). To establish the lemma, we need to show that ω is not optimal. To do so, we construct $\tilde{\omega}_i, \tilde{\omega}_j \geqslant 0$ such that $\tilde{\omega}_i + \tilde{\omega}_j = \omega_i + \omega_j$ and $(\tilde{\omega}_i, \tilde{\omega}_j) \neq (\omega_i, \omega_j)$, with $\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j \succcurlyeq_{\text{icv}} \omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j$, where $\succcurlyeq_{\text{icv}}$ is the increasing concave order (a.k.a second-order stochastic dominance). Because $h : \mathbb{R} \to \mathbb{R}$ given by $h(z) := -\mathbb{E}[(b - z - \sum_{k\in N\setminus\{i,j\}} \omega_k\boldsymbol{\varphi}_k)_+]$ is (weakly) increasing and concave, finding such $\tilde{\omega}_i, \tilde{\omega}_j$ would show that ω is not the unique minimizer of $\hat{\omega} \mapsto \mathbb{E}[(b - \hat{\omega}\cdot\boldsymbol{\varphi})_+]$—the objective in condition (i) of Theorem 1—and so is not optimal.

Now, define $\gamma := \frac{\alpha(\omega_i+\omega_j)}{\alpha^2\omega_i+\omega_j} > 0$, and let $\tilde{\omega}_i := \gamma\frac{1}{\alpha}\omega_j$ and $\tilde{\omega}_j := \gamma\alpha\omega_i$. By construction, $\tilde{\omega}_i + \tilde{\omega}_j = \omega_i + \omega_j$. Moreover, $(\tilde{\omega}_i, \tilde{\omega}_j) \neq (\omega_i, \omega_j)$—obviously if $\omega_i = 0 < \tilde{\omega}_i$, and otherwise because $\frac{\tilde{\omega}_j}{\tilde{\omega}_i} = \alpha\frac{\alpha\omega_i}{\omega_j} < \alpha < \frac{\omega_j}{\omega_i}$. It thus remains to show that

$\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j \succcurlyeq_{\text{icv}} \omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j$. To that end, first observe that

$$\begin{aligned}
\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j &= \gamma\tfrac{\omega_j}{\alpha}\boldsymbol{\varphi}_i + \gamma\alpha\omega_i\boldsymbol{\varphi}_j \\
&= \gamma\tfrac{\omega_j}{\alpha}\boldsymbol{\varphi}_i + \gamma\omega_i(\alpha\boldsymbol{\varphi}_j + \beta) - \gamma\omega_i\beta \\
&\succcurlyeq_{\text{icv}} \gamma\omega_i\boldsymbol{\varphi}_i + \gamma\tfrac{\omega_j}{\alpha}(\alpha\boldsymbol{\varphi}_j + \beta) - \gamma\omega_i\beta \\
&= \gamma(\omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j) + \gamma\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta \\
&= \gamma\left[\omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j)\right] + \gamma\left[\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta + \omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right],
\end{aligned}$$

where the inequality comes from Theorem 4.A.37 of Shaked and Shanthikumar (2007). Next, we establish that

$$\begin{aligned}
&\gamma\left[\omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j)\right] + \gamma\left[\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta + \omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right] \\
\succcurlyeq_{\text{icv}}\ &\omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j) + \gamma\left[\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta + \omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right]
\end{aligned}$$

To do so, observe that $\gamma = 1 - \frac{(1-\alpha)}{(\alpha^2\omega_i+\omega_j)}(\omega_j - \alpha\omega_i) \leqslant 1$ and $\mathbf{z} := \omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j)$ has zero mean. Because a constant shift obviously preserves $\succcurlyeq_{\text{icv}}$, we need only observe $\gamma\mathbf{z} \succcurlyeq_{\text{icv}} \mathbf{z}$, which follows directly from Jensen's inequality.[21]

Therefore,

$$\begin{aligned}
\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j &\succcurlyeq_{\text{icv}} \gamma\left[\omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j)\right] + \gamma\left[\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta + \omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right] \\
&\succcurlyeq_{\text{icv}} \omega_i(\boldsymbol{\varphi}_i - \bar{\theta}_i) + \omega_j(\boldsymbol{\varphi}_j - \bar{\theta}_j) + \gamma\left[\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta + \omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right] \\
&= \omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j + \gamma\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)\beta - (1-\gamma)\left(\omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right),
\end{aligned}$$

Because $\beta \geqslant (1-\alpha)\frac{\bar{\theta}_i+\alpha\bar{\theta}_j}{1+\alpha}$, it will therefore follow that $\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j \succcurlyeq_{\text{icv}} \omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j$ if we establish that

$$\lambda := \gamma\left(\tfrac{\omega_j}{\alpha} - \omega_i\right)(1-\alpha)\tfrac{\bar{\theta}_i+\alpha\bar{\theta}_j}{1+\alpha} - (1-\gamma)\left(\omega_i\bar{\theta}_i + \omega_j\bar{\theta}_j\right)$$

is nonnegative. And indeed, $\lambda = \frac{(1-\alpha)(\alpha\omega_i-\omega_j)^2}{(1+\alpha)(\alpha^2\omega_i+\omega_j)}(\bar{\theta}_i - \bar{\theta}_j)$, so the lemma will follow as long as we have $\bar{\theta}_i \geqslant \bar{\theta}_j$. For this ranking, note Theorem 1.B.42 of Shaked and Shanthikumar (2007) implies $\mathbb{E}[\boldsymbol{\varphi}_i] \geqslant \mathbb{E}\left[\alpha\boldsymbol{\varphi}_j + \beta\right]$, i.e.,

$$\bar{\theta}_i \geqslant \alpha\bar{\theta}_j + \beta \geqslant \alpha\bar{\theta}_j + (1-\alpha)\tfrac{\bar{\theta}_i+\alpha\bar{\theta}_j}{1+\alpha} = \bar{\theta}_i - \tfrac{2\alpha}{1+\alpha}(\bar{\theta}_i - \bar{\theta}_j).$$

Hence, $\bar{\theta}_i \geqslant \bar{\theta}_j$, as required. □

The following lemma sharpens the previous one by showing the weight ranking result often holds strictly. Whereas the previous lemma's proof uses the characterization of optimal weights as a minimax strategy, the following one uses the characterization as Minimizer's best response.

Lemma 4: *Suppose constants $\alpha \in (0,1]$ and $\beta \geqslant (1-\alpha)\frac{\bar{\theta}_i+\alpha\bar{\theta}_j}{1+\alpha}$ are such that*

$$\boldsymbol{\varphi}_i \succcurlyeq_{\text{rh}} \alpha\boldsymbol{\varphi}_j + \beta$$

and $\beta > 0$. Then, the optimal weight vector ω cannot satisfy $\alpha\omega_i = \omega_j > 0$.

[21] For $\eta : \mathbb{R} \to \mathbb{R}$ concave, $\mathbb{E}\eta(\gamma\mathbf{z}) \geqslant \mathbb{E}\left[\gamma\eta(\mathbf{z}) + (1-\gamma)\eta(0)\right] = \gamma\mathbb{E}\eta(\mathbf{z}) + (1-\gamma)\eta(\mathbb{E}\mathbf{z}) \geqslant \mathbb{E}\eta(\mathbf{z})$.

Proof. Consider any $\omega \in \Delta N$ with $\alpha\omega_i = \omega_j > 0$, with a view to showing it cannot be optimal. Defining the random variables $\tilde{\boldsymbol{\varphi}}_j := \alpha\boldsymbol{\varphi}_j + \beta$ and $\mathbf{y} := \frac{1}{\omega_i}\left(b - \sum_{k\in N\setminus\{i,j\}} \omega_k \boldsymbol{\varphi}_k\right) + \beta$, observe that $x_\omega(\boldsymbol{\theta}) = \mathbb{1}_{\boldsymbol{\varphi}_i + \tilde{\boldsymbol{\varphi}}_j \leqslant \mathbf{y}}$. Meanwhile, the random variables $\boldsymbol{\varphi}_i, \tilde{\boldsymbol{\varphi}}_j, \mathbf{y}$ are independent of $\mathbf{y}$ and $\boldsymbol{\varphi}_i \geqslant_{\text{rh}} \tilde{\boldsymbol{\varphi}}_j$.

Now, let us observe that $\mathbb{E}[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] \geqslant \mathbb{E}[\tilde{\boldsymbol{\varphi}}_j \mid \omega \cdot \boldsymbol{\varphi} \leqslant b]$. Indeed, this inequality is equivalent to showing $\mathbb{E}\eta(\tilde{\boldsymbol{\varphi}}_j, \boldsymbol{\varphi}_i) \geqslant 0$, where $\eta : \mathbb{R}^2 \to \mathbb{R}$ is given by $\eta(s,t) := (t-s)\mathbb{E}[\mathbb{1}_{s+t\leqslant \mathbf{y}}]$. Because $\eta(s,t) + \eta(t,s) = 0$ for every $s,t \in \mathbb{R}$ and η is nonincreasing in its first argument (as a product of two nonnegative nonincreasing functions) on $\{(s,t) \in \mathbb{R}^2 : \ s \leqslant t\}$, the inequality follows directly from Theorem 1.B.48 of Shaked and Shanthikumar (2007).

Hence, ω satisfies

$$\begin{aligned}\mathbb{E}[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] &\geqslant \mathbb{E}[\tilde{\boldsymbol{\varphi}}_j \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] \\ &= \alpha\mathbb{E}[\boldsymbol{\varphi}_j \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] + \beta.\end{aligned}$$

Assume now, for a contradiction, that ω is optimal. In this case, Theorem 1(ii) yields

$$\mathbb{E}[\boldsymbol{\varphi}_i \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] = \mathbb{E}[\boldsymbol{\varphi}_j \mid \omega \cdot \boldsymbol{\varphi} \leqslant b] =: \hat{\theta}.$$

Because the (interior-probability) event that $\omega\cdot\boldsymbol{\varphi} \leqslant b$ is the event that the bounded random variable $\boldsymbol{\varphi}_i$ [resp. $\boldsymbol{\varphi}_j$] lies below some random variable independent of it, it follows that $\hat{\theta} < \mathbb{E}[\boldsymbol{\varphi}_i] = \bar{\theta}_i$ [resp. $\hat{\theta} < \bar{\theta}_j$]. Therefore, $\frac{\bar{\theta}_i + \alpha\bar{\theta}_j}{1+\alpha} > \hat{\theta}$, implying $\beta > (1-\alpha)\hat{\theta}$. Hence, $\hat{\theta} \geqslant \alpha\hat{\theta} + \beta > \hat{\theta}$, a contradiction. □

We now reproduce the statement of Theorem 2.

Theorem (Ranking allocation weights)**:** *If $\boldsymbol{\varphi}_i \geqslant_{\text{rh}} \boldsymbol{\varphi}_j + \beta$ for some $\beta \geqslant 0$, then the optimal vector of allocation weights ω satisfies $\omega_i \geqslant \omega_j$. Moreover, $\omega_i > \omega_j$ whenever $\beta > 0$ and $\omega_j > 0$.*

Proof of Theorem 2. The first statement is exactly Lemma 3, specialized to the case of $\alpha = 1$. Given this result, the second statement corresponds exactly to Lemma 4, specialized to the case of $\alpha = 1$. □

A.3. Proofs for Section 5

We now reproduce the statement of Proposition 1.

Proposition (Optimal posted price is unanimous)**:** *Some unanimous posted-price mechanism is optimal among IC and IR collective posted-price mechanisms.*

Proof of Proposition 1. Consider an arbitrary collective posted-price mechanism (x, m) with price p. Let us show a unanimous posted price performs better.[22]

If $p \geqslant b$, then the profit associated with the mechanism is always nonpositive, and so a unanimous posted price with price in $(\max_{i\in N} \underline{\theta}_i, b)$ is more profitable.

[22] Our proof establishes any IC and IR collective posted price that is not almost surely identical to a unanimous one is *strictly* worse than some unanimous posted price.

Now, suppose $p < b$. For any agent $i \in N$ and $\theta_i \in (p, \bar{\theta}_i]$, IR implies $X_i^x(\theta_i) = 0$—and so $x(\theta_i, \boldsymbol{\theta}_{-i})$ must be zero almost surely. It follows that $x(\boldsymbol{\theta}) \leqslant x^U(\boldsymbol{\theta})$ almost surely, where x^U is the allocation rule

$$x^U(\theta) := \mathbb{1}_{\theta_j \leqslant p \ \forall j \in N}$$

associated with a unanimous posted price of p. Hence, $(b-p)\mathbb{E}[x(\boldsymbol{\theta})] \leqslant (b-p)\mathbb{E}[x^U(\boldsymbol{\theta})]$—strictly so unless $x(\boldsymbol{\theta}) = x^U(\boldsymbol{\theta})$ almost surely. Therefore, the unanimous posted-price mechanism (x^U, px^U) yields a higher profit.

Having shown every collective posted price is outperformed by some unanimous posted price, it remains to note that an optimal posted price exists. Any posted price outside of $(\max_{i \in N} \underline{\theta}_i, b)$ yields a nonpositive profit, whereas unanimous posted prices in this interval yield strictly positive profit. It thus suffices to show the buyer has some preferred price in $[\max_{i \in N} \underline{\theta}_i, b]$—which follows from compactness of this interval and continuity of the objective $p \mapsto (b-p)\prod_{i \in N} F_i(p)$. $\square$

Lemma 5: *If x is an optimal allocation rule, then $X_i^x(\cdot)$ is continuous on $(\underline{\theta}_i, \bar{\theta}_i)$ for every $i \in N$ with $\omega_i < 1$, and is nonconstant on $(\underline{\theta}_i, \bar{\theta}_i)$ if the optimal weights ω have $\omega_i > 0$.*

Proof. Let $\omega \in \Delta N$ be optimal, and let $X_i := X_i^{x_\omega}$ for each $i \in N$. Essential uniqueness of the optimal allocation rule (assured by Theorem 1) means it suffices to show $X_i(\cdot)$ is continuous for every $i \in N$, and is nonconstant on $(\underline{\theta}_i, \bar{\theta}_i)$ if the optimal weights ω have $\omega_i > 0$.[23]

First, let us see any given $i \in \text{supp}(\omega)$ is nonconstant. Indeed, a nonempty open neighborhood in Θ_{-i} exists such that $\omega \cdot (\underline{\theta}_i, \theta_{-i}) < b < \omega \cdot \varphi(\bar{\theta}_i, \theta_{-i})$ for any θ_{-i} in this neighborhood.[24] Because $\boldsymbol{\theta}_{-i}$ has full support and x is decreasing, it follows that $\lim_{\theta_i \searrow \underline{\theta}_i} X_i(\theta_i) < \lim_{\theta_i \nearrow \bar{\theta}_i} X_i(\theta_i)$. Hence, X_i is not constant on $(\underline{\theta}_i, \bar{\theta}_i)$.

Next, let us show that any $i \in N$ with $\omega_i < 1$ has X_i continuous. For each $\theta_i \in \Theta_i$, the interim probability of trade is given by

$$X_i^x(\theta_i) = \mathbb{P}\left\{ b - \sum\nolimits_{j \in N \setminus \{i\}} \omega_j \varphi_j(\boldsymbol{\theta}_j) \leqslant \omega_i \varphi_i(\theta_i) \right\}.$$

Recall that $\{\boldsymbol{\theta}_j\}_{j \in N}$ are independent and atomlessly distributed, ω_{-i} is nonzero, and $\varphi_i(\cdot)$ is continuous. It follows that the random variable on the left side of the above inequality is atomlessly distributed, while the quantity on the right side varies continuously with θ_i. Hence, X_i^x is continuous, as desired. $\square$

We now reproduce the statement of Proposition 2.

Proposition (Posted prices are suboptimal)**:** *If at least two $j \in N$ have $b < \bar{\theta}_j$, then no collective posted-price mechanism is optimal.*

[23] Essential uniqueness implies $X_i^x(\boldsymbol{\theta}_i) = X_i(\boldsymbol{\theta}_i)$ almost surely. Because $\boldsymbol{\theta}_i$ has convex support and X_i^x is monotone, it then follows (after establishing continuity of X_i) that the two functions are identical on $(\underline{\theta}_i, \bar{\theta}_i)$.

[24] Indeed, these inequalities hold whenever all of $\{\varphi_j(\theta_j)\}_{j \in N \setminus \{i\}}$ are within ϵ of b, where $\epsilon > 0$ is smaller than $\omega_i \min\{b - \underline{\theta}_i,\ \varphi_i(\bar{\theta}_i) - b\}$. This condition describes an open neighborhood because $\{\varphi_j\}_{j \in N}$ are continuous.

Proof of Proposition 2. Given Proposition 1, we need only show the unanimous posted-price mechanism is not an optimal mechanism for any price. Let ω denote the optimal weight vector assured by Theorem 1, fix some $i \in N$ such that $\omega_i > 0$, and let X_i denote i's interim allocation rule induced by a unanimous posted-price mechanism. By iterated expectations, constants p and $\bar{\mathtt{x}}$ exist such that every $\theta_i \in \Theta_i$ has $X_i(\theta_i) = \bar{\mathtt{x}}\mathbb{1}_{\theta_i \leq p}$. The function X_i therefore cannot be both continuous and nonconstant on $(\underline{\theta}_i, \bar{\theta}_i)$—it is discontinuous there if $\underline{\theta}_i < p < \bar{\theta}_i$ and $\bar{\mathtt{x}} \neq 0$, and is constant there otherwise. Given that the last assertion of Theorem 1 tells us $\omega_i < 1$, Lemma 5 thus delivers the proposition. □

ONLINE APPENDIX

B. Supporting analysis for Section 6

The following lemma shows that the share-weighted average of N independent types has a well-behaved distribution if each component does, and documents features of this distribution at the edges of its support.[25]

LEMMA 6: *Let G denote the cumulative distribution function of $\mathbf{v} = \sigma \cdot \boldsymbol{\theta}$.*

- *The distribution G admits a continuous density g which is strictly positive on the interior of its support $[\sigma \cdot \underline{\theta}, \sigma \cdot \bar{\theta}]$.*
- *As $v \nearrow \sigma \cdot \bar{\theta}$, we have*

$$\frac{g(v)}{(\sigma \cdot \bar{\theta} - v)^N} \to \frac{1}{(N-1)!} \prod_{i \in N} \frac{f_i(\bar{\theta}_i)}{\sigma_i}.$$

- *As $v \searrow \sigma \cdot \underline{\theta}$, we have*

$$\frac{g(v)}{(v - \sigma \cdot \underline{\theta})^{N-1}} \to \frac{1}{(N-1)!} \prod_{i \in N} \frac{f_i(\underline{\theta}_i)}{\sigma_i}.$$

- *As $v \searrow \sigma \cdot \underline{\theta}$, we have*

$$\frac{G(v)}{(v - \sigma \cdot \underline{\theta})^N} \to \frac{1}{N} \cdot \frac{1}{(N-1)!} \prod_{i \in N} \frac{f_i(\underline{\theta}_i)}{\sigma_i} \quad \textit{and} \quad \frac{\Phi(v) - v}{v - \sigma \cdot \underline{\theta}} \to \frac{1}{N},$$

where $\Phi(v) := v + \frac{G(v)}{g(v)}$.

Proof. For each $n \in N$, let G_n denote the CDF of $\sum_{i=1}^n \sigma_i \boldsymbol{\theta}_i$. The support of G_n is $\left[\sum_{i=1}^n \sigma_i \underline{\theta}_i,\ \sum_{i=1}^n \sigma_i \bar{\theta}_i\right]$, and when $n > 1$ any z in this support has

$$G_n(z) \;=\; \int_{\underline{\theta}_n}^{\bar{\theta}_n} G_{n-1}(z - \sigma_n \theta_n) f_n(\theta_n)\, \mathrm{d}\theta_n.$$

Because $G_1(v_1) = F_1(\frac{v_1}{\sigma_1})$ for every v_1, it follows that G_1 is continuously differentiable on its support with derivative $g_1(v_1) = \frac{1}{\sigma_1} f_1(\frac{v_1}{\sigma_1})$. Then, by induction on n, every $n \in N$ has G_n continuously differentiable on the interior of its support with the associated density at z in its support given by

$$g_n(z) = \int_{\underline{\theta}_n}^{\bar{\theta}_n} g_{n-1}(z - \sigma_n \theta_n) f_n(\theta_n)\, \mathrm{d}\theta_n.$$

[25] Whereas previous analyses apply readily to the case in which f_i may fail to be continuous and strictly positive at the endpoints of its support (like Example 1), the analysis of this section makes use of the fact that $\lim_{\theta_i \searrow \underline{\theta}_i} f_i(\underline{\theta}_i)$ and $\lim_{\theta_i \nearrow \bar{\theta}_i} f_i(\bar{\theta}_i)$ are both in $(0, \infty)$. Nevertheless, our qualitative results can be adapted to the case of power distributions—with $F_i(\theta_i) = \theta_i^\alpha$ for $\alpha > 0$—albeit with the threshold $\frac{2}{N+1}$ being replaced with the threshold $\frac{\alpha+1}{N\alpha+1}$.

Also by induction, g_n is strictly positive on the interior of its support because f_n is and (in the case of $n > 1$) g_{n-1} is. This establishes the first bullet.

To see the fourth bullet would follow from the third, note L'Hôpital's rule yields

$$\lim_{v \searrow \underline{\theta}_1} \tfrac{G(v)}{(v-\underline{\theta}_1)^N} = \lim_{v \searrow \underline{\theta}_1} \tfrac{g(v)}{N(v-\underline{\theta}_1)^{N-1}},$$

and note $\frac{\Phi(v)-v}{v-\underline{\theta}_1} = \frac{G(v)}{(v-\underline{\theta}_1)g(v)}$. So it remains to show the second and third bullets. Because the two are identical up to relabeling, we prove only the second bullet.

For any $\varepsilon > 0$ and any $n \in N$, let

$$h_n(\varepsilon) := (n-1)! \left[\prod_{i=1}^{n} \tfrac{\sigma_i}{f_i(\underline{\theta}_i)}\right] g_n\left(\varepsilon + \sum_{i=1}^{n} \sigma_i \underline{\theta}_i\right).$$

We need to show that $\frac{h_N(\varepsilon)}{\varepsilon^{N-1}} \to 1$ as $\varepsilon \searrow 0$. Let us show by induction that every $n \in N$ has $\frac{h_n(\varepsilon)}{\varepsilon^{n-1}} \to 1$ as $\varepsilon \searrow 0$, which will then deliver the lemma. For the base case, note that

$$\tfrac{h_1(\varepsilon)}{\varepsilon^0} = \tfrac{\sigma_1}{f_1(\underline{\theta}_1)} g_1\left(\varepsilon + \sigma_1 \underline{\theta}_1\right) = \tfrac{f_1\left(\underline{\theta}_1 + \frac{\varepsilon}{\sigma_1}\right)}{f_1(\underline{\theta}_1)},$$

which converges to 1 as $\varepsilon \searrow 0$.

For the inductive step, suppose $n > 1$ and that the limit equation holds for $n-1$. Then, any small enough $\varepsilon > 0$ has

$$\begin{aligned}
h_n(\varepsilon) &= (n-1)\tfrac{\sigma_n}{f_n(\underline{\theta}_n)} \int_{\underline{\theta}_n}^{\bar{\theta}_n} h_{n-1}\left(\varepsilon - \sigma_n\left[\theta_n - \underline{\theta}_n\right]\right) f_n(\theta_n)\, \mathrm{d}\theta_n \\
&= \tfrac{n-1}{f_n(\underline{\theta}_n)} \int_{\underline{\theta}_n}^{\underline{\theta}_n + \frac{\varepsilon}{\sigma_n}} h_{n-1}\left(\varepsilon - \sigma_n\left[\theta_n - \underline{\theta}_n\right]\right) f_n(\theta_n)\sigma_i\, \mathrm{d}\theta_n \\
&= \tfrac{n-1}{f_n(\underline{\theta}_n)} \int_0^{\varepsilon} h_{n-1}(\tilde{\varepsilon}) f_n(\underline{\theta}_n + \tfrac{\varepsilon-\tilde{\varepsilon}}{\sigma_n})\, \mathrm{d}\tilde{\varepsilon} \\
\implies \tfrac{h_n(\varepsilon)}{\varepsilon^{n-1}} - 1 &= \tfrac{h_n(\varepsilon)}{\varepsilon^{n-1}} - \tfrac{1}{\varepsilon^{n-1}} \int_0^{\varepsilon} (n-1)\tilde{\varepsilon}^{n-2}\, \mathrm{d}\tilde{\varepsilon} \\
&= (n-1)\tfrac{1}{\varepsilon} \int_0^{\varepsilon} \left(\tfrac{\tilde{\varepsilon}}{\varepsilon}\right)^{n-2} \left[\tfrac{h_{n-1}(\tilde{\varepsilon})}{\tilde{\varepsilon}^{n-2}} \tfrac{f_n\left(\underline{\theta}_n + \frac{\varepsilon-\tilde{\varepsilon}}{\sigma_n}\right)}{f_n(\underline{\theta}_n)} - 1\right] \mathrm{d}\tilde{\varepsilon},
\end{aligned}$$

which converges to zero (because the integrand does uniformly) as $\varepsilon \searrow 0$, as required. □

The next lemma shows how virtual costs from our group setting and the single-agent analogue can be ranked for very high and very low types.

Lemma 7: *Suppose $\{F_i\}_{i \in N}$ all coincide (so $\omega = (\frac{1}{N}, \ldots, \frac{1}{N})$ is optimal). Let G, g, and Φ be as defined in Lemma 6. Then:*

- *Every $v \in \Theta_1 \backslash \{\bar{\theta}_1\}$ close enough to $\bar{\theta}_1$ has $\Phi(v) > \varphi_1(\bar{\theta}_1)$.*
- *If $\sigma_i < \frac{2}{N+1}$ for every $i \in N$, then every $\theta \in \Theta \backslash \{\underline{\theta}\}$ close enough to $\underline{\theta}$ has $\omega \cdot \varphi(\theta) > \Phi(\sigma \cdot \theta)$.*

- *If $\sigma_i > \frac{2}{N+1}$ for some $i \in N$, then some $\eta \in \mathbb{R}^N_{++}$ exists such that every sufficiently small $\varepsilon > 0$ has $\omega \cdot \varphi(\underline{\theta} + \varepsilon\eta) < \Phi(\sigma \cdot (\underline{\theta} + \varepsilon\eta))$.*

Proof. All three parts follow from Lemma 6. First, as $v \nearrow \bar{\theta}_1$, that lemma tells us $g(v) \to 0$ so that $\Phi(v) \to \infty$. Meanwhile, that f_1 is continuous and strictly positive implies φ_1 is bounded. Hence, large enough $v \in [\underline{\theta}_1, \bar{\theta}_1)$ have $\Phi(v) > \varphi_1(\bar{\theta}_1)$.

Toward the second and third bullets, let us write $o(\theta - \underline{\theta})$ for any function of $\theta \in \Theta$ with $\frac{o(\theta-\underline{\theta})}{\|\theta-\underline{\theta}\|} \xrightarrow{\theta \searrow \underline{\theta}} 0$.[26] Lemma 6 tells us $\lim_{v \searrow \underline{\theta}_1} \frac{\Phi(v)-v}{v-\underline{\theta}_1} = \frac{1}{N}$, so that

$$\Phi(\sigma \cdot \theta) - \underline{\theta}_1 = \sigma \cdot \theta - \underline{\theta}_1 + \tfrac{1}{N}(\sigma \cdot \theta - \underline{\theta}_1) + o(\theta - \underline{\theta}) = \tfrac{N+1}{N}\sigma \cdot (\theta - \underline{\theta}) + o(\theta - \underline{\theta}).$$

Meanwhile, as $\theta_1 \searrow \underline{\theta}_1$, both $f_1(\theta_1)$ and $\frac{F_1(\theta_1)}{\theta_1 - \underline{\theta}_1}$ converge to $f_1(\underline{\theta}_1)$, so that

$$\tfrac{\varphi_1(\theta_1)-\underline{\theta}_1}{\theta_1-\underline{\theta}_1} = 1 + \tfrac{\varphi_1(\theta_1)-\theta_1}{\theta_1-\underline{\theta}_1} = 1 + \tfrac{F_1(\theta_1)}{(\theta_1-\underline{\theta}_1)f_1(\theta_1)} \to 2.$$

So $\varphi_i(\theta_i) - \underline{\theta}_i = 2(\theta_i - \underline{\theta}_i) + o(\theta - \underline{\theta})$, implying $\omega \cdot \varphi(\theta) - \underline{\theta}_1 = 2\omega \cdot (\theta - \underline{\theta}) + o(\theta - \underline{\theta})$. Therefore,

$$\begin{aligned} \omega \cdot \varphi(\theta) - \Phi(\sigma \cdot \theta) &= \left(2\omega - \tfrac{N+1}{N}\sigma\right) \cdot (\theta - \underline{\theta}) + o(\theta - \underline{\theta}) \\ &= \tfrac{N+1}{N}\left(\tfrac{2}{N+1}\mathbf{1}_N - \sigma\right) \cdot (\theta - \underline{\theta}) + o(\theta - \underline{\theta}). \end{aligned}$$

We now pursue the second bullet. If $\sigma_i < \frac{2}{N+1}$ for every $i \in N$, then the vector $\frac{N+1}{N}\left(\frac{2}{N+1}\mathbf{1}_N - \sigma\right)$ has strictly positive entries, so that $\omega \cdot \varphi(\theta) - \Phi(\sigma \cdot \theta) > 0$ for sufficiently small $\theta \in \Theta \backslash \{\underline{\theta}\}$.

Finally, to establish the third bullet, suppose some $i \in N$ has $\sigma_i > \frac{2}{N+1}$. Then, some $\gamma \in (0,1)$ exists such that

$$\gamma\left(\tfrac{2}{N+1} - \sigma_i\right) + (1-\gamma)\max_{j \in N}\left(\tfrac{2}{N+1} - \sigma_j\right) < 0.$$

Then $\eta \in \mathbb{R}^N_{++}$ with $\eta_i = \gamma$ and every other entry equal to $\frac{1-\gamma}{N-1}$ is as desired. □

Now, we introduce a notion of (utilitarian) efficiency ranking of allocation rules.

Definition 5: *Given an allocation rule x, the **surplus** generated by x in state $\theta \in \Theta$ is*

$$s_x(\theta) := x(\theta)(b - \sigma \cdot \theta).$$

*Given two allocation rules x and $\tilde{x}$, say x is **ex-ante more efficient** than $\tilde{x}$ if*

$$\mathbb{E}\left[s_x(\boldsymbol{\theta})\right] > \mathbb{E}\left[s_{\tilde{x}}(\boldsymbol{\theta})\right];$$

*and say x is **ex-post more efficient** than $\tilde{x}$ if*

$$\mathbb{P}\left\{s_x(\boldsymbol{\theta}) \geqslant s_{\tilde{x}}(\boldsymbol{\theta})\right\} = 1 \text{ and } \mathbb{P}\left\{s_x(\boldsymbol{\theta}) > s_{\tilde{x}}(\boldsymbol{\theta})\right\} > 0.$$

The next definition initializes language to discuss incentive properties and optimality of mechanisms in the single-agent benchmark.

[26] Note this property is independent of the norm because any two norms on $\mathbb{R}^N$ have bounded ratio.

Definition 6: *Say a mechanism* (x, m) *is* ***single-agent incentive compatible (SIC)*** *if*

$$\theta \in \text{argmax}_{\hat{\theta}\in\Theta}\left[m(\hat{\theta}) - \sigma \cdot \theta x(\hat{\theta})\right], \ \forall \theta \in \Theta,$$

that is, report $\hat{\theta} = \theta$ *maximizes the expected payoff of type profile* θ *over all possible reports in* Θ*. Say the mechanism is* ***single-agent individually rational (SIR)*** *if*

$$m(\theta) - \sigma \cdot \theta x(\theta) \geqslant 0, \ \forall \theta \in \Theta,$$

that is, the expected payoff of type profile θ*, when reporting truthfully, is nonnegative. A* ***single-agent-optimal mechanism*** *is an SIC and SIR mechanism that generates weakly higher buyer profit than any other SIC and SIR mechanism.*[27] *A* ***single-agent-optimal allocation rule*** *is any allocation rule* x *such that* (x, m) *is a single-agent-optimal mechanism for some* m*.*

Say an allocation rule x *is* ***single-agent implementable*** *if some transfer rule* m *exists such that the mechanism* (x, m) *is SIR; and say* x *is* ***single-agent monotone*** *if*

$$x(\sigma \cdot \theta) \leqslant x(\sigma \cdot \tilde{\theta}), \ \forall \theta, \tilde{\theta} \in \Theta \text{ with } \sigma \cdot \theta > \sigma \cdot \tilde{\theta}.$$

The next lemma shows that any single-agent-optimal mechanism is bounded between two cutoff mechanisms for the aggregated cost, in which the cutoffs solve a first-order condition equating the benefit of trade to single agent's virtual cost.

Lemma 8: *Some smallest and largest* $\underline{p}_b, \bar{p}_b \in (\sigma \cdot \underline{\theta}, \ \sigma \cdot \bar{\theta})$ *exist such that* $\Phi(\underline{p}_b) = \Phi(\bar{p}_b) = b$*, where* Φ *is as defined in Lemma 6. Moreover, some single-agent-optimal allocation rule exists, and any single-agent-optimal allocation rule* x *satisfies* $\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant\underline{p}_b} \leqslant x(\boldsymbol{\theta}) \leqslant \mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant\bar{p}_b}$ *almost surely.*

Proof. Lemma 6 tells us Φ is continuous on $(\sigma \cdot \underline{\theta}, \ \sigma \cdot \bar{\theta})$, and that $\Phi(v)$ converges to $\sigma \cdot \underline{\theta}$ [resp. ∞] as $v \searrow \sigma \cdot \underline{\theta}$ [resp. $v \nearrow \sigma \cdot \bar{\theta}$]. Therefore, the set $\left\{p \in (\sigma \cdot \underline{\theta}, \ \sigma \cdot \bar{\theta}) : \ \Phi(p) = b\right\}$ is closed and bounded away from $\{\sigma \cdot \underline{\theta}, \ \sigma \cdot \bar{\theta}\}$ (hence compact), the set is nonempty by the intermediate value theorem, and every price $p \in (\sigma \cdot \underline{\theta}, \ \sigma \cdot \bar{\theta})$ strictly below [resp. above] this set has $\Phi(p) < b$ [resp. $\Phi(p) > b$]. In particular, this set of prices has a smallest and largest element, $\underline{p}_b$ and $\bar{p}_b$, respectively.

Now define the allocation rule x^* by

$$x^*(\theta) := \mathbb{1}_{\sigma\cdot\theta\leqslant\underline{p}_b} + x(\theta)\mathbb{1}_{\sigma\cdot\theta\in(\underline{p}_b,\bar{p}_b]}.$$

Because x is $[0, 1]$-valued, a given $\theta \in \Theta$ has $\mathbb{1}_{\sigma\cdot\theta\leqslant\underline{p}_b} \leqslant x(\theta) \leqslant \mathbb{1}_{\sigma\cdot\theta\leqslant\bar{p}_b}$ if and only if $x(\theta) = x^*(\theta)$. It therefore remains to show $x(\boldsymbol{\theta}) = x^*(\boldsymbol{\theta})$ almost surely.

[27] Note, in this this short-hand, a single-agent-optimal mechanism/allocation is optimal *for the buyer* in the single-agent setting, and is not the preferred mechanism of the agent.

To show this equality, note that (straightforwardly adapting standard results from unidimensional mechanism design) a given allocation rule $\tilde{x}$ is single-agent implementable if and only if it is single-agent monotone, and that the maximum buyer value attainable by an SIC and SIR mechanism with allocation rule $\tilde{x}$ is $\mathbb{E}\{\tilde{x}(\boldsymbol{\theta})[b - \Phi(\sigma \cdot \boldsymbol{\theta})]\}$. Existence of an optimal allocation rule then follows from the observation that $\tilde{x} \mapsto \mathbb{E}\{\tilde{x}(\boldsymbol{\theta})[b - \Phi(\sigma \cdot \boldsymbol{\theta})]\}$ is a weak*-continuous function on the weak*-compact set $\tilde{\mathcal{X}}$.

Now, by construction (and since $\underline{p}_b \leqslant \bar{p}_b$), the allocation rule x^* is single-agent monotone because x is. Therefore, single-agent optimality of x tells us

$$\begin{aligned} 0 &\geqslant \mathbb{E}\{x^*(\boldsymbol{\theta})[b - \Phi(\sigma \cdot \boldsymbol{\theta})]\} - \mathbb{E}\{x(\boldsymbol{\theta})[b - \Phi(\sigma \cdot \boldsymbol{\theta})]\} \\ &= \mathbb{E}\left\{[1 - x(\boldsymbol{\theta})][b - \Phi(\sigma \cdot \boldsymbol{\theta})]\,\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} \leqslant \underline{p}_b} + x(\boldsymbol{\theta})[\Phi(\sigma \cdot \boldsymbol{\theta}) - b]\,\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} > \bar{p}_b}\right\}. \end{aligned}$$

Because a nonnegative random variable can have nonpositive expectation only if said random variable is almost surely zero, it follows that the random variable $[1 - x(\boldsymbol{\theta})][b - \Phi(\sigma \cdot \boldsymbol{\theta})]\,\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} \leqslant \underline{p}_b} + x(\boldsymbol{\theta})[\Phi(\sigma \cdot \boldsymbol{\theta}) - b]\,\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} > \bar{p}_b}$ is almost surely zero. Equivalently, $[1 - x(\boldsymbol{\theta})]\,\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} \leqslant \underline{p}_b} + x(\boldsymbol{\theta})\mathbb{1}_{\sigma \cdot \boldsymbol{\theta} > \bar{p}_b}$ is almost surely zero. Thus, $x(\boldsymbol{\theta}) = x^*(\boldsymbol{\theta})$ almost surely, as required. □

Finally, we prove an efficiency ranking result

Proposition 3: *Suppose $\{F_i\}_{i \in N}$ all coincide.*

(i) If b is large enough, then any optimal allocation rule for our group setting is ex-post more efficient than any single-agent-optimal allocation rule.

(ii) If b is small enough, then any optimal allocation rule for our group setting is ex-ante less efficient than any single-agent-optimal allocation rule.

Moreover, this efficiency ranking is an ex-post ranking if $\sigma_i < \frac{2}{N+1}$ for every $i \in N$ (in particular, if σ is close enough to ω), and is not an ex-post ranking if $\sigma_i > \frac{2}{N+1}$ for some $i \in N$ (in particular if σ is close enough to δ_i).

Proof. Because $\{F_i\}_{i \in N}$ all coincide, it follows from the uniqueness part of Theorem 1 that every allocation rule in our model agrees almost everywhere with x_ω, where $\omega = (\frac{1}{N}, \ldots, \frac{1}{N})$. Moreover, because ex-ante and ex-post efficiency rankings are both invariant to probability-zero changes to an allocation rule, we can prove the result by comparing single-agent-optimal allocation rules to x_ω. In what follows, let Φ be as defined in Lemma 6; let $\underline{p}_b, \bar{p}_b$ be as defined in Lemma 8; and use the notation $\mathbf{y} >_{\mathbb{P}} \tilde{\mathbf{y}}$ to say that the random variables $\mathbf{y}$ and $\tilde{\mathbf{y}}$ have $\mathbf{y} \geqslant \tilde{\mathbf{y}}$ almost surely with $\mathbb{P}\{\mathbf{y} > \tilde{\mathbf{y}}\} > 0$.

First, let us show x_ω is ex-post more efficient than single-agent-optimal allocation rules when b is high enough. To that end, note Lemma 7 tells us every $v \in [\underline{\theta}_1, \bar{\theta}_1)$ in some neighborhood of $\bar{\theta}_1$ in Θ has $\Phi(v) > \varphi_1(\bar{\theta}_1)$. Because $\theta \mapsto \omega \cdot \varphi(\theta)$ is continuous and strictly increasing, some $b^* \in (\bar{\theta}_1, \varphi_1(\bar{\theta}_1))$ exists such that every $\theta \in \Theta \backslash \{\bar{\theta}\}$ with $\omega \cdot \varphi(\theta) > b^*$ is in said neighborhood. Now, take any $b \in [b^*, \varphi_1(\bar{\theta}_1))$ and any single-agent monotone allocation rule x; we want to show x is ex-post less efficient than x_ω. To see it, given any $\theta \in \Theta \backslash \{\bar{\theta}\}$

with $\omega \cdot \varphi(\theta) \geqslant b$, note that any $\tilde{\theta} \in \Theta$ with $\tilde{\theta} \geqslant \theta$ has $\omega \cdot \varphi(\tilde{\theta}) \geqslant b$ and so $\Phi(\sigma \cdot \tilde{\theta}) > \varphi_1(\bar{\theta}_1) > b$. Thus, any $\theta \in \Theta$ with $\omega \cdot \varphi(\theta) \geqslant b$ has $\sigma \cdot \theta > \bar{p}_b$, where $\bar{p}_b$ is as given by Lemma 8. Because φ is continuous, it follows that any $\theta \in \Theta$ with $\omega \cdot \varphi(\theta)$ close enough to b also has $\sigma \cdot \theta > \bar{p}_b$. Therefore, $x_\omega(\boldsymbol{\theta}) >_{\mathbb{P}} \mathbb{1}_{\sigma \cdot \boldsymbol{\theta} \leqslant \bar{p}_b}$. Lemma 8 then implies $x_\omega(\boldsymbol{\theta}) >_{\mathbb{P}} x(\boldsymbol{\theta})$. Finally, because $b > \bar{\theta}_1$, it follows that $s_{x_\omega}(\boldsymbol{\theta}) >_{\mathbb{P}} s_x(\boldsymbol{\theta})$. That is, x_ω is ex-post more efficient than x.

Next, specializing to the case in which each $i \in N$ has $\sigma_i < \frac{2}{N+1}$, let us show x_ω is ex-post less efficient than single-agent-optimal allocation rules when b is low enough. To that end, note Lemma 7 tells us every $\theta \in \Theta \backslash \{\underline{\theta}\}$ in some neighborhood of $\underline{\theta}$ has $\omega \cdot \varphi(\theta) > \Phi(\sigma \cdot \theta)$; let $b_* \in (\underline{\theta}_1, \bar{\theta}_1)$ be small enough that every $\theta \in \Theta \backslash \{\underline{\theta}\}$ with $\omega \cdot \theta \leqslant b_*$ or $\sigma \cdot \theta \leqslant b_*$ is in said neighborhood. Now, take any $b \in (\underline{\theta}_1, b_*]$ and any single-agent-optimal allocation rule x; we want to show x is ex-post less efficient than x_ω. To see it, given any $\theta \in \Theta \backslash \{\underline{\theta}\}$ with $\omega \cdot \varphi(\theta) \leqslant b$, note that any $\tilde{\theta} \in \Theta$ with $\tilde{\theta} \leqslant \theta$ has $\omega \cdot \tilde{\theta} \leqslant \omega \cdot \varphi(\tilde{\theta}) \leqslant b$ and so $\Phi(\sigma \cdot \tilde{\theta}) < \omega \cdot \varphi(\tilde{\theta}) \leqslant b$. Thus, any $\theta \in \Theta$ with $\omega \cdot \varphi(\theta) \leqslant b$ has $\omega \cdot \theta < \underline{p}_b$, where $\underline{p}_b$ is as given by Lemma 8. Because φ is continuous, it follows that any $\theta \in \Theta$ with $\omega \cdot \varphi(\theta)$ close enough to b also has $\omega \cdot \theta < \underline{p}_b$. Therefore, $x_\omega(\boldsymbol{\theta}) <_{\mathbb{P}} \mathbb{1}_{\omega \cdot \boldsymbol{\theta} \geqslant \underline{p}_b}$. Lemma 8 then implies $x_\omega(\boldsymbol{\theta}) <_{\mathbb{P}} x(\boldsymbol{\theta})$. Moreover, because $\underline{p}_b < \Phi(\underline{p}_b) = b$, it follows from Lemma 8 that, almost surely, either $x(\boldsymbol{\theta}) = 0$ or $\sigma \cdot \boldsymbol{\theta} < b$. It follows that $s_{x_\omega}(\boldsymbol{\theta}) <_{\mathbb{P}} s_x(\boldsymbol{\theta})$. That is, x is ex-post more efficient than x_ω.

Now, specializing to the case in which some $i \in N$ has $\sigma_i > \frac{2}{N+1}$, let us show x_ω is *not* ex-post more efficient than single-agent-optimal allocation rules when b is low enough. To that end, let note Lemma 7 delivers some $\eta \in \mathbb{R}^N_{++}$ such that $\omega \cdot \varphi(\underline{\theta} + \varepsilon\eta) < \Phi\left(\sigma \cdot (\underline{\theta} + \varepsilon\eta)\right)$ for all sufficiently small $\varepsilon > 0$. Let $\theta(\varepsilon) := \underline{\theta} + \varepsilon\eta$ for every ε, and for any $b \in (\underline{\theta}_1, \varphi_1(\bar{\theta}_1))$, let $\varepsilon_b := \frac{\bar{p}_b - \underline{\theta}_1}{\sigma \cdot \eta}$ so that $\sigma \cdot \theta(\varepsilon_b) = \bar{p}_b$. That $b = \Phi(\bar{p}_b) > \bar{p}_b$ then implies $\sigma \cdot \theta(\varepsilon_b) < b$ and $\bar{p}_b \searrow \underline{\theta}_1$ as b does, and so too does ε_b. Therefore, whenever $b \in (\underline{\theta}_1, \varphi_1(\bar{\theta}_1))$ is sufficiently small, we have that $\theta(\varepsilon_b)$ is interior in Θ and

$$\omega \cdot \varphi(\theta(\varepsilon_b)) < \Phi(\sigma \cdot \theta(\varepsilon_b)) = \Phi(\bar{p}_b) = b.$$

Let us fix such a small b and any single-agent allocation rule x, with a view to showing x is not ex-post more efficient than x_ω. Because φ is continuous, then, some $\hat{\theta} > \theta(\varepsilon_b)$ in the interior of Θ is close enough to $\theta(\varepsilon_b)$ to ensure $\omega \cdot \varphi(\hat{\theta}) < b$ and $\sigma \cdot \hat{\theta} < b$. Thus,

$$\bar{p}_b < \sigma \cdot \hat{\theta} < b \text{ and } \omega \cdot \varphi(\hat{\theta}) < b.$$

Again by continuity, every θ in some neighborhood of $\hat{\theta}$ satisfies the same three inequalities. Therefore, $\mathbb{P}\left\{\bar{p}_b < \sigma \cdot \boldsymbol{\theta} < b \text{ and } \omega \cdot \varphi(\boldsymbol{\theta}) < b\right\} > 0$. Lemma 8 then tells us $\mathbb{P}\left\{x(\boldsymbol{\theta}) = 0,\ x_\omega(\boldsymbol{\theta}) = 1, \text{ and } \omega \cdot \varphi(\boldsymbol{\theta}) < b\right\} > 0$. Thus, x is not ex-post more efficient than x_ω.

Finally, returning to the case of general σ, let us show x_ω is ex-ante less efficient than single-agent-optimal allocation rules when b is low enough. For any $b \in (\underline{\theta}_1, \varphi_1(\bar{\theta}_1))$, let $x_{\sigma,b}$ denote some single-agent-optimal allocation rule, and let

$$S(\sigma, b) := \mathbb{E}\left\{x_{\sigma,b}(\boldsymbol{\theta})(b - \sigma \cdot \boldsymbol{\theta})\right\}$$

denote the surplus it generates. We want to show that

$$S(\sigma, b) > \mathbb{E}\left\{x_\omega(\boldsymbol{\theta})(b - \sigma \cdot \boldsymbol{\theta})\right\}$$

when b is close enough to $\underline{\theta}_1$. Because $\{\boldsymbol{\theta}_i\}_{i\in N}$ are i.i.d., we know that $\mathbb{E}\left\{x_\omega(\boldsymbol{\theta})(b - \boldsymbol{\theta}_i)\right\}$ is the same for each $i \in N$, so that $\mathbb{E}\left\{x_\omega(\boldsymbol{\theta})(b - \sigma \cdot \boldsymbol{\theta})\right\} = \mathbb{E}\left\{x_\omega(\boldsymbol{\theta})(b - \omega \cdot \boldsymbol{\theta})\right\}$. Meanwhile, the ex-post efficiency ranking of this proof's second paragraph implies (by taking expectations) the ex-ante efficiency ranking $S(\omega, b) > \mathbb{E}\left\{x_\omega(\boldsymbol{\theta})(b - \omega \cdot \boldsymbol{\theta})\right\}$ for all sufficiently small b. The proposition will therefore follow if we can show $S(\sigma, b) \geqslant S(\omega, b)$ when $b \in (\underline{\theta}_1, \varphi_1(\bar{\theta}_1))$ is close enough to $\underline{\theta}_1$. We now pursue this ranking.

Let $\gamma := \frac{f_1(\underline{\theta}_1)^N}{(N-1)!}\prod_{i\in N}\frac{1}{\sigma_i} > 0$. In what follows, we use Lemma 6's calculations of the behavior of G, g, and Φ around $\underline{\theta}_1$. First, for $p \in (\underline{\theta}_1, \bar{\theta}_1)$, we have

$$\begin{aligned}
\tfrac{1}{(p-\underline{\theta}_1)^{N+1}}\mathbb{E}\left[\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant p}\left(\sigma \cdot \boldsymbol{\theta} - \underline{\theta}_1\right)\right] &= \tfrac{1}{(p-\underline{\theta}_1)^{N+1}}\int_{\underline{\theta}_1}^{p}(v - \bar{\theta}_1)g(v)\ \mathrm{d}v \\
&= \tfrac{\gamma}{(p-\underline{\theta}_1)^{N+1}}\int_{\underline{\theta}_1}^{p}(v - \bar{\theta}_1)^N\ \mathrm{d}v \\
&\quad + \tfrac{1}{p-\underline{\theta}_1}\int_{\underline{\theta}_1}^{p}\left(\tfrac{v-\underline{\theta}_1}{p-\underline{\theta}_1}\right)^N\left[\tfrac{g(v)}{(v-\underline{\theta}_1)^{N-1}} - \gamma\right]\ \mathrm{d}v \\
&= \tfrac{\gamma}{N+1} + \tfrac{1}{p-\underline{\theta}_1}\int_{\underline{\theta}_1}^{p}\left(\tfrac{v-\underline{\theta}_1}{p-\underline{\theta}_1}\right)^N\left[\tfrac{g(v)}{(v-\underline{\theta}_1)^{N-1}} - \gamma\right]\ \mathrm{d}v \\
&\xrightarrow{p\searrow\underline{\theta}_1} \tfrac{\gamma}{N+1}.
\end{aligned}$$

Moreover, we have

$$\tfrac{\Phi(p)-\underline{\theta}_1}{p-\underline{\theta}_1} = 1 + \tfrac{\Phi(p)-p}{p-\underline{\theta}_1} \xrightarrow{p\searrow\underline{\theta}_1} 1 + \tfrac{1}{N} = \tfrac{N+1}{N}.$$

Therefore,

$$\begin{aligned}
&\tfrac{1}{[\Phi(p)-\underline{\theta}_1]^{N+1}}\mathbb{E}\left\{\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant p}\left[\Phi(p) - \sigma \cdot \boldsymbol{\theta}\right]\right\} \\
= &\left[\tfrac{p-\underline{\theta}_1}{\Phi(p)-\underline{\theta}_1}\right]^N \tfrac{G(p)}{(p-\underline{\theta}_1)^N} - \left[\tfrac{p-\underline{\theta}_1}{\Phi(p)-\underline{\theta}_1}\right]^{N+1} \tfrac{1}{(p-\underline{\theta}_1)^{N+1}}\mathbb{E}\left[\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant p}\left(\sigma \cdot \boldsymbol{\theta} - \underline{\theta}_1\right)\right] \\
\xrightarrow{p\searrow\underline{\theta}_1} &\left(\tfrac{N}{N+1}\right)^N \tfrac{\gamma}{N} - \left(\tfrac{N}{N+1}\right)^{N+1} \tfrac{\gamma}{N+1} \\
= &\tfrac{N^{N-1}}{(N+1)^{N+2}(2N+1)}\gamma.
\end{aligned}$$

Meanwhile, any $b \in (\underline{\theta}_1, \bar{\theta}_1)$ has $\underline{\theta}_1 < \underline{p}_b \leqslant \bar{p}_b < \Phi(\bar{p}_b) = b$, so that $\underline{p}_b, \bar{p}_b \searrow \underline{\theta}_1$ as b does. We can therefore specialize the above calculation to deduce

$$\tfrac{1}{(b-\underline{\theta}_1)^{N+1}}\mathbb{E}\left[\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant \underline{p}_b}\left(b - \sigma \cdot \boldsymbol{\theta}\right)\right] \text{ and } \tfrac{1}{(b-\underline{\theta}_1)^{N+1}}\mathbb{E}\left[\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant \bar{p}_b}\left(b - \sigma \cdot \boldsymbol{\theta}\right)\right]$$

both converge to $\frac{N^{N-1}}{(N+1)^{N+2}(2N+1)}\gamma$ as $b \searrow \underline{\theta}_1$. Now, because every $v \leqslant \bar{p}$ has $v \leqslant \Phi(\bar{p}) = b$, Lemma 8 implies

$$\mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant \underline{p}_b}\left(b - \sigma \cdot \boldsymbol{\theta}\right) \leqslant x_{\sigma,b}(\boldsymbol{\theta})\left(b - \sigma \cdot \boldsymbol{\theta}\right) \leqslant \mathbb{1}_{\sigma\cdot\boldsymbol{\theta}\leqslant \bar{p}_b}\left(b - \sigma \cdot \boldsymbol{\theta}\right),$$

so that

$$\begin{aligned}\tfrac{1}{(b-\underline{\theta}_1)^{N+1}}S(\sigma,b) &\xrightarrow{b\searrow\underline{\theta}_1} \tfrac{N^{N-1}}{(N+1)^{N+2}(2N+1)}\gamma \\ &= \tfrac{N^{N-1}f_1(\underline{\theta}_1)^N}{(N-1)!(N+1)^{N+2}(2N+1)}\prod_{i\in N}\frac{1}{\sigma_i} \\ &= \frac{\tilde{\gamma}}{\prod_{i\in N}\sigma_i}, \text{ where}\end{aligned}$$

where $\tilde{\gamma} := \tfrac{N^{N-1}f_1(\underline{\theta}_1)^N}{(N-1)!(N+1)^{N+2}(2N+1)} > 0$. Note that this calculation specializes to

$$\tfrac{1}{(b-\underline{\theta}_1)^{N+1}}S(\omega,b) \xrightarrow{b\searrow\underline{\theta}_1} \frac{\tilde{\gamma}}{\left(\frac{1}{N}\right)^N}.$$

Therefore,

$$\frac{S(\sigma,b)}{S(\omega,b)} \xrightarrow{b\searrow\underline{\theta}_1} \frac{\left(\frac{1}{N}\right)^N}{\prod_{i\in N}\sigma_i} = \left[\frac{\frac{1}{N}\sum_{i\in N}\sigma_i}{\left(\prod_{i\in N}\sigma_i\right)^{\frac{1}{N}}}\right]^N.$$

The inequality of arithmetic and geometric means (AM-GM) tells us that this limit ratio is strictly greater than 1 if $\sigma \neq \omega$, so that $S(\sigma,b) \leqslant S(\omega,b)$ when $b \in (\underline{\theta}_1, \bar{\theta}_1)$ is sufficiently small. The proposition follows. □

C. Supporting analysis for Section 7

C.1. Dominant strategies

In light of the revelation principle, we formalize more demanding incentive constraints through direct mechanisms below.

Definition 7: *Say a mechanism* (x,m) *is* ***dominant-strategy incentive compatible (DIC)*** *if*

$$\theta_i \in \operatorname{argmax}_{\hat{\theta}_i\in\Theta_i}\left\{m(\hat{\theta}_i,\theta_{-i}) - \theta_i x(\hat{\theta}_i,\theta_{-i})\right\},\ \forall i\in N,\ \forall\theta\in\Theta; \qquad \text{(DIC)}$$

A mechanism is DIC if an agent finds truthful reporting dominant in the direct revelation game; that is, he would willingly report truthfully even if he knew others' reported types.

We showed in Lemma 1 that for a given allocation rule, interim montonicity is equivalent to BIC implementability. Said differently, we showed that being able to BIC-implement an allocation rule with agent-specific transfers is equivalent to being able to do so with only collective transfers. Moreover, Theorem 1 explicitly characterizes the allocation rule from optimal BIC and IR mechanisms, showing it stipulates trade if and only if the benefit to the buyer exceeds the player-weighted virtual cost. Notice, though, that this allocation rule is monotone in the agents' *profile* of types. If our seller could engage in agent-specific transfers, such monotonicity would render the same allocation rule DIC implementable too. Therefore, a natural conjecture is that (as in single-good auction settings) our seller can attain DIC at no additional cost.

The following result shows the above natural conjecture is false: the restriction to DIC mechanisms is with loss of optimality for the seller. Optimal mechanisms must leverage agents' uncertainty about others' realized types.

Proposition 4 (Dominance binds)**:** *If at least two $j \in N$ have $b < \bar{\theta}_j$, then no DIC mechanism is optimal.*

The proof of Proposition 4 leverages the fact that the essentially unique optimal allocation rule is bang-bang—every type profile leads to a deterministic trade outcome. The main thrust of our proof is a structural lemma that characterizes the full class of DIC bang-bang mechanisms, as summarized in two properties. The first property concerns the transfer: It can be decomposed into a price (p) that will be paid if and only if trade occurs and a subsidy that will be paid to the sellers whether or not trade occurs. The second property gives a representation of the allocation rule: trade is determined by the price and $\mathcal{J}$, a collection of subsets of N such that the good is sold if and only if, for some $J \in \mathcal{J}$, every agent in J agrees to the purchase at price p.

The proof of the structural lemma proceeds in two steps. First, we show the transfer rule is constant among type profiles leading to certain trade, and constant among type profiles leading to non-trade, which leads directly to the price/subsidy form. To prove this property, consider any two type profiles θ and θ' such that $x(\theta) = x(\theta')$; say this trade probability is equal to 1, the alternative case being analogous. Letting θ^* be a type profile that is coordinatewise higher than both θ and θ', we construct a finite sequence of type profiles such that the first type profile in the sequence is θ and the last is θ^*, the type profiles get coordinatewise higher as the sequence progresses, and consecutive entries in the sequence differ in only one agent's type. But then, because DIC (for the agent whose type is raised in a given increment of the sequence) implies x must be monotone, it follows that every type profile in the sequence generates probability 1 of trade. Hence, DIC (again, for the agent whose type is incremented) implies consecutive sequence members yield an identical transfer. A symmetric argument applies to θ', so that $m(\theta') = m(\theta^*) = m(\theta)$. Hence, any DIC-implementing transfer takes the given price-subsidy form. The second property that the structural lemma establishes is the structure on the allocation rule. Given that the mechanism is incentive-equivalent to a collective posted price of p, DIC implies (fixing a realization of others' types) the trade decision must be identical for all types of agent i below p and for all types of agent i above p. Hence, the allocation rule is essentially a decreasing $\{0,1\}$-valued transformation of the vector-valued function $\theta \mapsto (\mathbb{1}_{\theta_j \geqslant p})_{j \in N}$. The "coalitional" property amounts to a more explicit description of such functions.

C.2. Ex-post participation

Let us formulate a notion of ex-post individual rationality. As usual, we do so for direct mechanisms—for convenience and without loss. Say a mechanism (x, m) is

ex-post indiviually rational (epIR) if $m(\theta) - \theta_i x(\theta) \geqslant 0$ for every $\theta \in \Theta$ and $i \in N$.

The following lemma reduces IC-and-epIR implementability to the study of the allocation rule and agents' interim values.

Lemma 9: *Given allocation rule x and $\underline{U} \in \mathbb{R}^N$, the following are equivalent:*

(i) Some transfer rule m exists such that the mechanism (x, m) is IC and epIR and gives interim utility $\underline{U}_i$ to type $\bar{\theta}_i$ of each agent i.

(ii) The allocation rule x is interim monotone, the quantities $\{\underline{U}_i + \mathbb{E}[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i]\}_{i\in N}$ all coincide, and every agent $i \in N$ and type $\theta_i \in \Theta_i$ have (letting $X_i := X_i^x$):

$$\int_{\theta_i}^{\bar{\theta}_i} X_i(\tilde{\theta}_i)\, \mathrm{d}\tilde{\theta}_i \geqslant \mathbb{E}\left[x(\theta_i, \boldsymbol{\theta}_{-i}) \left(\max_{j\in N\setminus\{i\}} \boldsymbol{\theta}_j - \theta_i\right)_+\right] - \underline{U}_i.$$

The above lemma shows how an epIR constraint can be formulated directly over allocation rules. The conditions given in the lemma amount to saying that, when the interim transfer rules are solved out from the allocation rule via the sellers' IC constraint and revenue equivalence, seller i's interim transfer is at least her interim expectation of the minimum transfer required to stop all sellers from walking away. This condition is trivially necessary, but we constructively show it to be sufficient too.

As a demonstration that the above characterization is useful, let us apply it to derive a sufficient condition for epIR to be without loss of optimality for our buyer.

Proposition 5 (Sufficient condition for epIR)**:** *Suppose $N = 2$ and $F_1 = F_2$. Then, some optimal mechanism is epIR if the virtual cost $\boldsymbol{\varphi}_1$ admits a nonincreasing density on its support.*

In particular, this proposition applies to the special case of Example 1 in which $\alpha_1 = \alpha_2 \leqslant 1$.

C.3. Pareto-optimal Mechanisms

Recall, a **Pareto-optimal mechanism**, is an IC and IR mechanism such that no alternative IC and IR mechanism delivers a weakly higher buyer profit, and a weakly higher agent i value for each agent i, with at least one of these $N + 1$ inequalities strict. Then, a **Pareto-optimal allocation** is any allocation rule x such that (x, m) is a Pareto-optimal mechanism for some m. In this subsection, we provide a characterization of which mechanisms are Pareto optimal, and explain the reasoning behind it.

Following standard arguments, one can show that any Pareto optimal mechanism can be represented as a solution to a program maximizing a weighted sum of values of the $N + 1$ individuals (N sellers and the buyer), and—because increasing the transfer by a constant preserves all constraints—the Pareto weight on the buyer (normalized to 1) is at least as high as the sum of weights $\{\lambda_i\}_{i\in N}$ on the agents. Conversely, we observe that any interim monotone allocation rule that maximizes such a weighted sum is Pareto optimal.[28]

We can therefore solve a family of programs much like the buyer's problem (BP), but with modified objective, to trace out the entire Pareto frontier. Vectors λ of Pareto weights are paired with endogenous allocation weights ω to describe the following class of allocation rules.

Definition 8: *Let $\Delta(2N)$ denote the set of all (λ, ω) with $\lambda, \omega \in \mathbb{R}_+^N$ and*

$$\sum_{i \in N} (\lambda_i + \omega_i) = 1.$$

*For any such (λ, ω), let the (λ, ω)-**allocation rule**, denoted by $x_{\lambda,\omega}$, be given by*

$$x_{\lambda,\omega}(\theta) := \mathbb{1}_{\lambda \cdot \theta + \omega \cdot \varphi(\theta) \leqslant b}.$$

We now state our main characterization theorem of this section. It characterizes Pareto-optimal allocation rules as those that weigh the benefit of trade against a weighted average of its *actual and virtual* costs.

Theorem 3 (Pareto-optimal allocations)**:** *The (λ, ω)-allocation rule is Pareto optimal for any $(\lambda, \omega) \in \Delta(2N)$ satisfying the following two equivalent conditions:*

1. $\omega \in \operatorname{argmin}_{\tilde{\omega}:\ (\lambda, \tilde{\omega}) \in \Delta N} \mathbb{E}[(b - \lambda \cdot \boldsymbol{\theta} - \tilde{\omega} \cdot \boldsymbol{\varphi})_+]$.
2. $\operatorname{supp}(\omega) \subseteq \operatorname{argmax}_{i \in N} \mathbb{E}\left[\boldsymbol{\varphi}_i \mid \lambda \cdot \boldsymbol{\theta} + \omega \cdot \boldsymbol{\varphi} \leqslant b\right]$.

Moreover, every Pareto-optimal allocation rule is essentially of this form.[29]

Finally, analogous to Proposition 2, it is natural to explore whether some Pareto optimal mechanism can be implemented via posted prices. That is, one can ask whether the suboptimality of collective posted prices was an artefact of our focus on the buyer-optimal mechanisms. However, as we show in Proposition 7, no Pareto optimal mechanism can be implemented as a collective posted price mechanism (if trade is neither unambiguously efficient nor inefficient). Therefore, this intuitive class of simple mechanisms is strictly suboptimal regardless of whether one favors the buyer or the seller.

C.4. Pre-market trade of land shares

Consider a game that extends our model by adding a pre-market phase in which the agents trade their shares. The buyer then observes the agents' shares and chooses a profit-maximizing mechanism. We study two different versions of this game: one in which agents must be paid proportionally to their shares, as we have required throughout this paper, and one without this constraint. We show that in the first regime, agents do not benefit from trade, but in the second regime they do.

[28] The latter observation would be obvious if all weights were strictly positive. We show it holds in our setting even with some zero weights, because the optimizer is essentially unique.

[29] The proof also establishes that, if the (λ, ω)- and $(\lambda, \tilde{\omega})$-allocation rules are both Pareto optimal, then they essentially coincide.

We start with some notation and then define the game. Let $\Sigma = \{\sigma = (\sigma_1, \ldots, \sigma_N) \in (0,1)^N : \sum_{i \in N} \sigma_i = 1\}$ be the set of possible profiles of shares that the agents might have. Agents' initial shares $\sigma = (\sigma_1, \ldots, \sigma_N) \in \Sigma$ are fixed but unknown to the buyer. For the present analysis, we assume $\{\boldsymbol{\theta}_i\}_{i \in N}$ are i.i.d. with distribution $F_i = F_1$. A mechanism is a profile of functions $(x, m_1, \ldots, m_N)$ where $x : \Theta \to [0,1]$ is the allocation rule and $m_i : \Theta \to \mathbb{R}$ is the transfer rule of agent i. Notice that here we are considering a more general class of mechanisms than the one studied throughout this paper because here we impose no structure relating the transfers of different agents to each other. Play proceeds as follows.

1. Seller 1 proposes shares $\hat{\sigma} \in \Sigma$ and lump-sum net transfers $\hat{\tau} \in \mathbb{R}^N$ with $\mathbb{1} \cdot \tau = 0$, and then the other agents sequentially vote on whether to accept the proposal. Realized shares σ' and transfers τ are then equal to $\hat{\sigma}$ and $\hat{\tau}$ if all accept the proposal, and equal to σ and $\vec{0}$ if anyone rejects.[30]
2. The buyer observes σ' and chooses a mechanism (x, m).

   In the **discriminatory-pricing** regime, the buyer can choose any mechanism. In the **uniform-pricing** regime, the buyer can choose any uniform-pricing mechanism—that is a mechanism $(x, m_1, \ldots, m_N)$ in which agents are paid proportionally to their chosen shares, $m_i = \sigma_i' \sum_{j \in N} m_j$.
3. Each agent i privately learns his type $\boldsymbol{\theta}_i$ drawn independently from F, decides whether to participate in the mechanism, and if he participates, what type $\hat{\theta}_i$ to report.
4. The good is sold with probability $x(\hat{\theta})$ and each agent i is paid $m_i(\hat{\theta})$. The payoff of agent i is then $\tau_i + m_i(\hat{\theta}) - \sigma_i' \boldsymbol{\theta}_i x(\hat{\theta})$.[31]

Our solution concept, which we simply call equilibrium for brevity, is perfect Bayesian equilibrium in which:

- Players do not signal what they do not know—hence, play from stage 2 onward corresponds to the mechanism design problem with shares σ' and type distribution $\boldsymbol{\theta} \sim \bigotimes_{i \in N} F_1$;
- The buyer-optimal mechanism is offered, and sellers all participate and truthfully report their types, for any realized shares σ'.[32]

We say that the game is **pre-market non-manipulable** if some equilibrium exists in which the agents choose $\sigma' = \sigma$ in the first stage and the buyer attains her optimal value (among all IC and IR mechanisms for shares σ). The following result shows that a uniform-pricing regime generates such non-manipulability.

[30] The specific bargaining protocol is immaterial, though we fix one for concreteness. What matters for our analysis is that the realized shares σ' are set to maximize sellers' sum of payoffs.

[31] A more natural specification would be $\tau_i + m_i(\hat{\theta}) + \sigma_i' \boldsymbol{\theta}_i [1 - x(\hat{\theta})]$, which explicitly takes into account that seller i has value $\sigma_i' \boldsymbol{\theta}_i$ (which depends on σ') if he retains his land. Because we maintain i.i.d. types for the present analysis, though, the difference $\sum_{i \in N} \mathbb{E}[\sigma_i' \boldsymbol{\theta}_i]$ does not vary with σ', and so will not affect pre-market non-manipulability. We therefore maintain the payoff specification of our main model for ease of comparison.

[32] The latter feature simplifies the analysis, but is not necessary. If we removed this equilibrium refinement, but enriched the model to allow the buyer to pay sellers even when some sellers do not participate, our results would remain unchanged.

Proposition 6 (Uniform pricing avoids pre-market trade)**:** *Suppose sellers' types are i.i.d. Then, the uniform-pricing game is pre-market non-manipulable, but the discriminatory-pricing game need not be.*

The proof shows that sellers' total surplus is invariant to their shares under uniform pricing, and shows a numerical example (an instance of Example 1) in which the sellers increase their total surplus by making their shares more symmetric.

D. Proofs for Section C

D.1. Proofs for Section C.1

Lemma 10: *Suppose that (x, m) is a DIC mechanism and $\theta, \theta' \in \Theta$ have $x(\theta) = x(\theta') \in \{0, 1\}$. Then $m(\theta) = m(\theta')$.*

Proof. Define $\theta^* := \theta \vee \theta'$ if $x(\theta) = x(\theta') = 0$, and $\theta^* := \theta \wedge \theta'$ if $x(\theta) = x(\theta') = 1$. We will observe that $m(\theta) = m(\theta^*) = m(\theta')$; by symmetry, it suffices to show $m(\theta) = m(\theta^*)$. To show it, define the type profile

$$\theta^\ell := (\theta_i^* \mathbb{1}_{i \leqslant \ell} + \theta_i \mathbb{1}_{i > \ell})_{i \in N} \in \Theta \text{ for each } \ell \in \{0, \ldots, N\} = N \cup \{0\}.$$

Observe, either $\theta^0 \leqslant \cdots \leqslant \theta^N$ and $x(\theta^0) = 0$, or $\theta^0 \geqslant \cdots \geqslant \theta^N$ and $x(\theta^0) = 1$. In either case, because x is weakly decreasing (due to DIC) and can only take values in $[0, 1]$, it follows by induction that $x(\theta^0) = \cdots = x(\theta^N)$. For each $i \in N$, because θ^i and θ^{i-1} differ only in the i coordinate and $x(\theta^{i-1}) = x(\theta^i)$, it follows from DIC (for agent i) that $m(\theta^{i-1}) = m(\theta^i)$. Thus, $m(\theta) = m(\theta^0) = \cdots = m(\theta^N) = m(\theta^*)$, as desired. □

Definition 9: *Say a mechanism (x, m) or an allocation rule x is **bang-bang** if $x(\boldsymbol{\theta}) \in \{0, 1\}$ almost surely.*

Lemma 11: *Suppose (x, m) is a DIC bang-bang mechanism. Then, some $p, s \in \mathbb{R}$ and $\mathcal{J} \subseteq 2^N$ exist such that, almost surely:*

(i) $m(\boldsymbol{\theta}) = px(\boldsymbol{\theta}) + s$;

(ii) $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{\boldsymbol{\theta}_j \leqslant p\}}$.

Moreover, we may assume without loss that no two members of $\mathcal{J}$ are nested, and that $\underline{\theta}_j < p < \bar{\theta}_j$ for each $j \in \bigcup \mathcal{J}$.

Proof. Fix a DIC mechanism (x, m) such that $x(\boldsymbol{\theta})$ almost surely in $\{0, 1\}$. By Lemma 10, some constants $m^L, m^H \in \mathbb{R}$ exists such that $m(\theta) = m^L$ [resp. m^H] for every $\theta \in \Theta$ with $x(\theta) = 0$ [resp. 1]. So, defining $p := m^H - m^L \geqslant 0$ and letting $s := m^L$, we have $m(\boldsymbol{\theta}) = px(\boldsymbol{\theta}) + s$ whenever $x(\boldsymbol{\theta}) \in \{0, 1\}$, an almost sure event.

Now, modifying x on an a null set, and similarly modifying the transfer rule to maintain $m = px + s$, we may assume without loss that x is (statewise) $\{0, 1\}$-valued.[33] DIC of the modified mechanism follows from DIC of the original one.

[33]For instance, if x is almost-surely constant, we can modify it to be constant; and otherwise, we can replace x with $\theta \mapsto \mathbb{1}_{x(\theta) > 0}$.

Next, we show x has the desired structure. Given an agent $i \in N$ and type realization $\theta_i \in \Theta_i$, his payoff from a reported type profile of $\hat{\theta}$ is $(p - \theta_i)x(\hat{\theta}) - s$, which is strictly increasing [resp. decreasing] in $x(\hat{\theta})$ if $\theta_i < p$ [resp. $\theta_i > p$]. Hence, given $\theta_{-i} \in \Theta_{-i}$ DIC implies that one the following three possibilities holds: $x(\cdot, \theta_{-i}) = 1$ globally, $x(\cdot, \theta_{-i}) = 0$ globally, or $x(\theta_i, \theta_{-i}) = 1$ [resp. $x(\theta_i, \theta_{-i}) = 0$] for each $\theta_i \in \Theta_i$ with $\theta_i < p$ [resp. $\theta_i > p$]. Hence, letting $\tilde{\Theta} := \prod_{i \in N}[\Theta_i \backslash \{p\}]$, some $y : \{0,1\}^N \to \{0,1\}$ exists such that every $\theta \in \tilde{\Theta}$ has $x(\theta) = y\left((\mathbb{1}_{\theta_i \leqslant p})_{i \in N}\right)$. Moreover, we may assume without loss that y is constant in its i coordinate if $p \leqslant \underline{\theta}_i$ or $p \geqslant \bar{\theta}_i$ for $i \in N$. Then, monotonicity of x implies y is monotone too. If we let $\tilde{\mathcal{J}} := \{J \subseteq N : \ y(\mathbb{1}_J) = 1\}$, then, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{\tilde{J} \in \tilde{\mathcal{J}}} \bigcap_{j \in \tilde{J}} \{\boldsymbol{\theta}_j \geqslant p\}}$ almost surely.

Define $\hat{\mathcal{J}} := \left\{ \{j \in \tilde{J} : \ \underline{\theta}_j < p\} : \ \tilde{J} \in \tilde{\mathcal{J}} \text{ with } \bar{\theta}_j > p \ \forall j \in \tilde{J} \right\}$. Then, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{\hat{J} \in \hat{\mathcal{J}}} \bigcap_{j \in \hat{J}} \{\boldsymbol{\theta}_j \leqslant p\}}$ almost surely, and $\underline{\theta}_j < p < \bar{\theta}_j$ for each $j \in \bigcup \hat{\mathcal{J}}$. Finally, let $\mathcal{J} := \{J \in \hat{\mathcal{J}} : \ \nexists \hat{J} \in \hat{\mathcal{J}} \text{ with } \hat{J} \subsetneq J\}$. Then, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{\boldsymbol{\theta}_j \leqslant p\}}$ almost surely, $\underline{\theta}_j < p < \bar{\theta}_j$ for each $j \in \bigcup \mathcal{J}$, and no two members of $\mathcal{J}$ are nested. Thus, $(p, s, \mathcal{J})$ is as required. □

Proof of Proposition 4. Let x be any DIC-implementable allocation rule. First, let $(p, s, \mathcal{J})$ be as delivered by Lemma 11 (with $\mathcal{J}$ chosen so that the "moreover" part of the lemma holds).

Let us show x it cannot be optimal. First, if $\mathcal{J}$ is either $\varnothing$ or $\{\varnothing\}$, then $\mathbb{E}[x(\boldsymbol{\theta})] \in \{0,1\}$, and so Theorem 1 says (given that $\underline{\theta}_i < b < \varphi_i(\bar{\theta}_i)$ for each $i \in N$) that x is not optimal. Second, if $i \in J \in \mathcal{J}$, then X_i^x is discontinuous at $p \in (\underline{\theta}_i, \bar{\theta}_i)$, implying (by Lemma 5 and since the last assertion of Theorem 1 tells us ω is nontrivial) that x is not an optimal allocation rule. □

D.2. Proofs for Section C.2

Proof of Lemma 9. First, define the transfer rule $\underline{m}$ by letting $\underline{m}(\theta) := \max_{i \in N} \theta_i x(\theta)$. Note that a transfer rule m is such that (x, m) is epIR if and only if $m \geqslant \underline{m}$.

Now, for each agent i, let M_i^* be as defined in the proof of Lemma 1. As explained in that proof, given a transfer rule m, the mechanism (x, m) is IC and gives high-type utility $\underline{U}_i$ to each agent i if and only if x is interim monotone and $M_i^m = M_i^* + \underline{U}_i$ for each agent i. So condition (i) holds if and only if x is interim monotone and some transfer rule m exists such that $m \geqslant \underline{m}$ and $M_i^m = M_i^* + \underline{U}_i$ for each agent i. Observe, the last condition also implies that $\{\underline{U}_i + \mathbb{E}\left[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i\right]\}_{i \in N}$ all coincide because (as noted in the proof of Lemma 1) each $i \in N$ has $\mathbb{E}\left[M_i^*(\boldsymbol{\theta})\right] = \mathbb{E}\left[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i\right]$.

To prove the lemma, it therefore suffices to show the following: Given a profile $(M_i)_{i \in N}$ of interim transfer rules such that $\{\mathbb{E}\left[M_i(\boldsymbol{\theta}_i)\right]\}_{i \in N}$ all coincide, the following are equivalent:

- Some transfer rule $m \geqslant \underline{m}$ has $M_i^m = M_i$ for each agent i;
- Each agent i has $M_i \geqslant M_i^{\underline{m}}$.

To see this equivalence delivers the lemma, note that the the inequality $M_i(\theta_i) \geqslant$

$M_i^{\underline{m}}(\theta_i)$ rearranges to exactly the inequality in the lemma's statement.

The first bullet immediately implies the second, because integration is monotone. To pursue the converse, suppose the second bullet holds, that is, $M_i \geqslant M_i^{\underline{m}}$ for each agent i. Let $\mathtt{m} := \mathbb{E}\left[M_i(\boldsymbol{\theta}_i) - M_i^{\underline{m}}(\boldsymbol{\theta}_i)\right]$, which is the same nonnegative quantity for every agent i. If $\mathtt{m} > 0$, then the transfer rule m given by

$$m(\theta) := \underline{m}(\theta) + \mathtt{m}^{-(N-1)} \prod_{i \in N} \left[M_i(\theta_i) - M_i^{\underline{m}}(\theta_i)\right]$$

is as desired; and if $\mathtt{m} = 0$, then the transfer rule m given by

$$m(\theta) := \underline{m}(\theta) + \max_{i \in N} \left[M_i(\theta_i) - M_i^{\underline{m}}(\theta_i)\right]$$

is as desired. Indeed, in both cases, $m \geqslant \underline{m}$ by construction; in the $\mathtt{m} > 0$ case, $M_i^m = M_i$ because agents' types are independent; and in the $\mathtt{m} = 0$ case each $M_i^m = M_i$ because types are independent and each agent $j \neq i$ has $M_j^m(\boldsymbol{\theta}_j) = M_j^{\underline{m}}(\boldsymbol{\theta}_j)$ almost surely. □

The following lemma simplifies the characterization of the previous lemma to understand when epIR is without loss of optimality in our buyer's problem, in the two-agent symmetric case.

Lemma 12: *Suppose $N = 2$ and $F_1 = F_2$. Let $\bar{z} := \boldsymbol{\varphi}_1(\bar{\theta}_1)$, let G denote the CDF of $\boldsymbol{\varphi}_1$, let $\lambda := \varphi_1^{-1} : [\underline{\theta}_1, \bar{z}] \to \Theta_1$ extended to be constant above $\bar{z}$, and let $\hat{(\cdot)} : \mathbb{R} \to \mathbb{R}$ be given by $\hat{y} := 2b - y$ (the reflection across b).*

Then, some optimal mechanism is epIR if and only if every $z \in [\underline{\theta}_1, b]$ has

$$G(z)\left[\lambda(\hat{z}) - \lambda(z)\right] + \int_z^{\hat{z}} \left[\lambda(\hat{y}) - \lambda(y)\right] \, \mathrm{d}G(y) + \int_{\hat{z}}^{\bar{z}} G(\hat{y}) \, \mathrm{d}\lambda(y) \geqslant 0.$$

Proof. Let x denote the $(\frac{1}{2}, \frac{1}{2})$-weighted allocation rule. From Theorem 1, we know that an optimal mechanism exists with allocation rule x and IR binding for both agents. Letting $X_1 := X_1^x$, define the function $\eta : \Theta_1 \to \mathbb{R}$ via

$$\eta(\theta_1) := \int_{\theta_1}^{\bar{\theta}_1} X_1(\tilde{\theta}_1) \, \mathrm{d}\tilde{\theta}_1 - \mathbb{E}\left[x(\theta_1, \boldsymbol{\theta}_2)\,(\boldsymbol{\theta}_2 - \theta_1)_+\right]$$

for each $\theta_1 \in \Theta_1$. Given symmetry and given Lemma 9, we know some optimal mechanism is epIR if the function η is globally nonnegative. Conversely, because (given Theorem 1) any optimal mechanism has binding IR and has an allocation rule that agrees with x almost surely, and because (as will be clear from our analysis below) η is continuous, it follows that nonnegativity of the function η is also necessary for some optimal mechanism to be epIR. The lemma will then follow if we show the inequalities in the lemma's statement characterize global nonnegativity of η.

Observe now that $G, \lambda, \hat{(\cdot)}$ are all continuous and monotone, and λ is strictly increasing. To see when $\eta(\theta_1) \geqslant 0$ for every $\theta_1 \in \Theta_1$, we equivalently characterize when $\eta(\lambda(z)) \geqslant 0$ for every $z \in [\underline{\theta}_1, \bar{z}]$.

Now, let us compute η more explicitly. Any $z \in [\underline{\theta}_1, \bar{z}]$ has

$$x(\lambda(z), \boldsymbol{\theta}_2) = \mathbb{1}_{\frac{1}{2}\varphi_1(\lambda(z))+\frac{1}{2}\boldsymbol{\varphi}_2 \leqslant b} = \mathbb{1}_{\boldsymbol{\varphi}_2 \leqslant \hat{z}},$$

and so (extending λ to equal $\bar{\theta}_1$ above $\bar{z}$),

$$\begin{aligned} \eta(\lambda(z)) &= \int_{\lambda(z)}^{\lambda(\bar{z})} X_1(\tilde{\theta}_1)\ \mathrm{d}\tilde{\theta}_1 - \mathbb{E}\left\{\mathbb{1}_{\boldsymbol{\varphi}_2 \leqslant \hat{z}} \left[\lambda(\boldsymbol{\varphi}_2) - \lambda(z)\right] \mathbb{1}_{\lambda(\boldsymbol{\varphi}_2) > \lambda(z)}\right\} \\ &= \int_z^{\bar{z}} X_1(\lambda(y))\ \mathrm{d}\lambda(y) - \mathbb{E}\left\{\mathbb{1}_{z < \boldsymbol{\varphi}_2 \leqslant \hat{z}} \left[\lambda(\boldsymbol{\varphi}_2) - \lambda(z)\right]\right\} \\ &= \int_z^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) - \mathbb{1}_{z<\hat{z}} \int_z^{\hat{z}} \left[\lambda(y) - \lambda(z)\right]\ \mathrm{d}G(y). \end{aligned}$$

Now, observe $\mathbb{1}_{z<\hat{z}} = \mathbb{1}_{z<b}$. Thus, if $z \geqslant b$, we have $\eta(\lambda(z)) = \int_z^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) \geqslant 0$.

So let us focus on the remaining case of $z < b$. In this case, note that $\hat{(\cdot)}$ is a continuous decreasing bijection on $[z, \hat{z}]$, and so

$$\begin{aligned} \int_z^{\hat{z}} G(\hat{y})\ \mathrm{d}\lambda(y) &= \int_{\hat{z}}^{z} G(y)\ \mathrm{d}[\lambda \circ \hat{(\cdot)}](y) \\ &= \left[G(y)\lambda(\hat{y})\right]_{y=\hat{z}}^{z} - \int_{\hat{z}}^{z} \lambda(\hat{y})\ \mathrm{d}G(y) \\ &= \left[G(z)\lambda(\hat{z}) - G(\hat{z})\lambda(z)\right] + \int_z^{\hat{z}} \lambda(\hat{y})\ \mathrm{d}G(y). \end{aligned}$$

Therefore,

$$\begin{aligned} \eta(\lambda(z)) &= \int_z^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) - \int_z^{\hat{z}} \left[\lambda(y) - \lambda(z)\right]\ \mathrm{d}G(y) \\ &= \int_z^{\hat{z}} G(\hat{y})\ \mathrm{d}\lambda(y) + \int_{\hat{z}}^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) + \int_z^{\hat{z}} \lambda(z)\ \mathrm{d}G(y) - \int_z^{\hat{z}} \lambda(y)\ \mathrm{d}G(y) \\ &= \left\{\left[G(z)\lambda(\hat{z}) - G(\hat{z})\lambda(z)\right] + \int_z^{\hat{z}} \lambda(\hat{y})\ \mathrm{d}G(y)\right\} \\ &\quad + \int_{\hat{z}}^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) + \lambda(z)\left[G(\hat{z}) - G(z)\right] - \int_z^{\hat{z}} \lambda(y)\ \mathrm{d}G(y) \\ &= G(z)\left[\lambda(\hat{z}) - \lambda(z)\right] + \int_z^{\hat{z}} \left[\lambda(\hat{y}) - \lambda(y)\right]\ \mathrm{d}G(y) + \int_{\hat{z}}^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y). \end{aligned}$$

Thus, η is globally nonnegative if and only if the last expression is globally nonnegative for each $z \in [\underline{\theta}_1, b)$, as required. □

Proof of Proposition 5. Let $\bar{z}, G, \lambda, \hat{(\cdot)}$ be as defined in the previous lemma. In light of that lemma, we need to show each $z \in [\underline{\theta}_1, b)$ has

$$G(z)\left[\lambda(\hat{z}) - \lambda(z)\right] + \int_z^{\hat{z}} \left[\lambda(\hat{y}) - \lambda(y)\right]\ \mathrm{d}G(y) + \int_{\hat{z}}^{\bar{z}} G(\hat{y})\ \mathrm{d}\lambda(y) \geqslant 0.$$

By hypothesis G admits a nonincreasing density g on its support. In this case, any $z \in [\underline{\theta}_1, b)$ has

$$\begin{aligned}
& G(z)\left[\lambda(\hat{z})-\lambda(z)\right]+\int_z^{\hat{z}}\left[\lambda(\hat{y})-\lambda(y)\right]\,\mathrm{d}G(y)+\int_{\hat{z}}^{\bar{z}} G(\hat{y})\,\mathrm{d}\lambda(y) \\
\geqslant\ & \int_z^{\hat{z}}\left[\lambda(\hat{y})-\lambda(y)\right]\,\mathrm{d}G(y) \\
=\ & \int_z^{b}\left[\lambda(\hat{y})-\lambda(y)\right]g(y)\,\mathrm{d}y+\int_b^{\hat{z}}\left[\lambda(\hat{y})-\lambda(y)\right]g(y)\,\mathrm{d}y \\
=\ & \int_z^{b}\left[\lambda(\hat{y})-\lambda(y)\right]\left[g(y)-g(\hat{y})\right]\,\mathrm{d}y \\
\geqslant\ & 0,
\end{aligned}$$

as required. □

D.3. Proofs for Section C.3

In this section, we prove our characterization of the Pareto frontier. We also extend some previous results for buyer-optimal mechanisms to the entire Pareto frontier.

D.3.1. Proof of Theorem 3

To simplify our algebra in what follows, let $\vec{y}$ denote the vector $y\mathbb{1}_N \in \mathbb{R}^N$ for any scalar $y \in \mathbb{R}$.

Definition 10: *Let Λ denote the set of all vectors $\lambda \in \mathbb{R}^N_+$ such that $\vec{1} \cdot \lambda \leqslant 1$.*

*Given $\lambda \in \Lambda$, a λ-**compatible** vector is any ω such that $(\lambda, \omega) \in \Delta(2N)$.*

The following lemma studies a program in which an allocation rule is chosen to maximize a weighted sum of utilities, the monotonicity property required by IC is ignored, the payment formula is assumed, and the constant on the payment formula is chosen to make IR bind for some agent. To state the lemma, for any $x \in \mathcal{X}$ or $\tilde{\mathcal{X}}$, define the profit level

$$\pi(x) := \min_{i \in N} \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right].$$

In light of Lemma 1, if $x \in \mathcal{X}$ is implementable, this profit level is the highest one consistent with IC and IR mechanisms that use allocation rule x.

Lemma 13: *Given any $\lambda \in \Lambda$, a unique solution exists to program*

$$\max_{x \in \tilde{\mathcal{X}}} \left\{(1 - \vec{1} \cdot \lambda)\pi(x) + \lambda \cdot \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right]\right\}.$$

This solution is given by the (λ, ω)-allocation rule, where ω is any λ-compatible vector satisfying the following two equivalent conditions:

(i) $\omega \in \text{argmin}_{\tilde{\omega}:\ (\lambda,\tilde{\omega})\in\Delta N}\, \mathbb{E}[(b - \lambda\cdot\boldsymbol{\theta} - \tilde{\omega}\cdot\boldsymbol{\varphi})_+]$.

(ii) $\text{supp}(\omega) \subseteq \text{argmax}_{i\in N}\, \mathbb{E}\left[\boldsymbol{\varphi}_i \mid \lambda\cdot\boldsymbol{\theta} + \omega\cdot\boldsymbol{\varphi} \leqslant b\right]$.

Proof. Substituting the definition of $\pi(x)$ and rearranging, the program's objective can be rewritten as

$$\min_{i\in N} \mathbb{E}\left\{x(\boldsymbol{\theta})\left[b - \lambda\cdot\boldsymbol{\theta} - (1 - \vec{1}\cdot\lambda)\boldsymbol{\varphi}_i\right]\right\}.$$

We can therefore follow the proof of Lemma 2 by modifying the two-player zero-sum game. Specifically, have Minimizer choose from the altered strategy space $(1 - \vec{1}\cdot\lambda)\Delta N$ of λ-compatible vectors, and change the objective to

$$\mathcal{G}_\lambda(x,\omega) := \mathbb{E}\left[x(\boldsymbol{\theta})\,(b - \lambda\cdot\boldsymbol{\theta} - \omega\cdot\boldsymbol{\varphi})\right].$$

Following exactly the proof of Lemma 2, mutatis mutandis, delivers the result. $\square$

Motivated by the above lemma, we will say a vector ω is **λ-optimal** if it is λ-compatible and satisfies the numbered conditions in Lemma 13.

In what follows, let $Z \subseteq \mathbb{R}\times\mathbb{R}^N$ denote the set

$$Z = \left\{\left(\pi,\ \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}\right) :\ x \in \mathcal{X},\ \pi\in\mathbb{R},\ \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] \geqslant \pi\ \forall i \in N\right\},$$

which is the set of payoff vectors induced by all mechanisms when the payment formula and IR are imposed. Given any $\tilde{Z} \subseteq \mathbb{R}\times\mathbb{R}^N$, say a point (π, u) is **Pareto optimal in** $\tilde{Z}$ if $(\pi,u)\in\tilde{Z}$ and no $(\tilde{\pi},\tilde{u}) \in \tilde{Z}\backslash\{(\pi,u)\}$ exists with $(\tilde{\pi},\tilde{u}) \geqslant (\pi,u)$.

The following lemma establishes a useful technical property of the payoff set Z and its Pareto frontier.

Lemma 14: *Every $z \in Z$ admits some $\tilde{z} \geqslant z$ that is Pareto optimal in Z.*

Proof. We begin with useful preliminary claim: The set $\{z \in Z :\ z \geqslant \underline{z}\}$ is compact for any $\underline{z} \in \mathbb{R}\times\mathbb{R}^N$. To show this fact, write $\underline{z} = (\underline{\pi}, \underline{u})$. Letting $\bar{\pi} := \min_{i\in N}\mathbb{E}\left[(b - \boldsymbol{\varphi}_i)_+\right]$, note that no $x \in \mathcal{X}$ and $\pi > \bar{\pi}$ can satisfy $\mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] \geqslant \pi\ \forall i \in N$. Because $\tilde{\mathcal{X}}$ is weak* compact (by Banach Alaoglu), the set

$$\left\{\left(\pi,\ \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}\right) :\ x \in \tilde{\mathcal{X}},\ \pi\in[\underline{\pi},\bar{\pi}],\ \mathbb{E}\left[x(\boldsymbol{\theta})(b - \boldsymbol{\varphi}_i)\right] \geqslant \pi\ \forall i \in N\right\}$$

is a continuous image of a compact space. Therefore, $\{z \in Z :\ z \geqslant \underline{z}\}$ is the intersection of the closed set $\mathbb{R}\times(\underline{u} + \mathbb{R}^N_+)$ with a compact set, and so is compact.

With the compactness claim in hand, we now establish the lemma. View Z as a subset of $\mathbb{R}^{\{0,\ldots,N\}}$, and let $z^{-1} := z \in Z$. For each $\tilde{z}\in Z$, let $Z(\tilde{z}) := \{\hat{z} \in Z : \hat{z} \geqslant \tilde{z}\}$, a nonempty (containing $\tilde{z}$) and compact subset of Z. Recursively for each $j \in \{0,\ldots,N\}$, we can therefore take $z^j \in \text{argmax}_{\tilde{z}\in Z(z^{j-1})}\, \tilde{z}_j$. By construction, $z \leqslant z^0 \leqslant \cdots \leqslant z^N$. Let us observe $\tilde{z} := z^N$ is Pareto optimal in Z. To that end, let $\hat{z} \in Z(\tilde{z})$; we want to show $\hat{z} \leqslant \tilde{z}$. And indeed, every $j \in \{0,\ldots,N\}$ has $\hat{z} \in Z(z^{j-1})$, so that $\hat{z}_j \leqslant z^j_j \leqslant \tilde{z}_j$. Therefore, $\hat{z} = \tilde{z}$, delivering the lemma. $\square$

The following lemma links Pareto optimality in the value set Z to the cutoff rule form.

Lemma 15: *Take any $\pi^* \in \mathbb{R}$ and $x^* \in \mathcal{X}$, and let $u^* := \mathbb{E}\left[x^*(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}^*$. The vector (π^*, u^*) is Pareto optimal in Z if and only if some $\lambda \in \Lambda$ and λ-optimal ω exist such that: $x^*(\boldsymbol{\theta}) = x_{\lambda,\omega}(\boldsymbol{\theta})$ almost surely; and $\pi^* \leqslant \pi(x_{\lambda,\omega})$, with equality if $\omega \neq \vec{0}$.*

Proof. Let us prove the following three conditions are equivalent:

(a) Payoff vector (π^*, u^*) is Pareto optimal in Z.

(b) Some $\lambda \in \Lambda$ exists such that

$$\begin{aligned}(\pi^*, x^*) \in \quad & \operatorname{argmax}_{(\pi,x)\in\mathbb{R}\times\mathcal{X}} \quad \left\{\pi + \lambda \cdot \left(\mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}\right)\right\} \\ & \text{s.t.} \qquad\qquad \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\varphi}_i)\right] \geqslant \pi \ \forall i \in N.\end{aligned}$$

(c) Some $\lambda \in \Lambda$ and λ-optimal ω exist such that: $x^*(\boldsymbol{\theta}) = x_{\lambda,\omega}(\boldsymbol{\theta})$ almost surely; and $\pi^* \leqslant \pi(x_{\lambda,\omega})$, with equality if $\omega \neq \vec{0}$.

First, let us see that conditions (b) and (c) are equivalent. To that end, fix $\lambda \in \Lambda$, and consider the program in condition (b), which can be rewritten as

$$\max_{(\pi,x)\in\mathbb{R}\times\mathcal{X}} \left\{(1 - \vec{1} \cdot \lambda)\pi + \lambda \cdot \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right]\right\} \text{ s.t. } \pi \leqslant \pi(x).$$

For any given $x \in \mathcal{X}$, the optimization for π is trivial to solve. The objective is weakly increasing in π (because $\vec{1} \cdot \lambda \leqslant 1$), strictly so if $\vec{1} \cdot \lambda < 1$. Therefore, condition (b) is satisfied if and only if:

- $\pi^* \leqslant \pi(x^*)$, with equality if $\vec{1} \cdot \lambda < 1$;
- $x^* \in \operatorname{argmax}_{x\in\mathcal{X}} \left\{(1 - \vec{1} \cdot \lambda)\pi(x) + \lambda \cdot \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right]\right\}$.

The equivalence then follows directly from Lemma 13.

Now, we establish condition (b) implies condition (a). To that end, suppose condition (b) holds, and take any $(\pi, u) \in Z$ with $(\pi, u) \geqslant (\pi^*, u^*)$; we want to show $(\pi, u) = (\pi^*, u^*)$. First, by definition of Z, some allocation rule x exists such that $x = \mathbb{E}\left[x(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}$. Then, that $(\pi, u) \geqslant (\pi^*, u^*)$ implies—because $(\pi, u) \mapsto \pi + \lambda \cdot (u - \vec{\pi})$ is weakly increasing—that (π, x) is also an optimal solution to the program in condition (b). Hence, x is an optimal solution to the program in Lemma 13. The uniqueness part of Lemma 13 therefore tells us $x(\boldsymbol{\theta}) = x^*(\boldsymbol{\theta})$ almost surely. By revenue equivalence (Myerson, 1981; Myerson and Satterthwaite, 1983), then, $u - u^* = (\pi - \pi^*)\vec{1}$. Hence,

$$(\pi - \pi^*)(1, \ -\vec{1}) = (\pi, u) - (\pi^*, u^*) \geqslant 0,$$

implying $\pi - \pi^* = 0$, and so $u = u^*$, as required.

Finally, let us show condition (a) implies condition (b). Supposing (π^*, u^*) is Pareto optimal in Z, we want to show some $\lambda \in \Lambda$ exists such that $(\pi^*, u^*) \in \operatorname{argmax}_{(\pi,u)\in Z} [\pi + \lambda \cdot u]$. First note, Z is the linear image of a set defined by linear inequality constraints on a convex domain; hence it is convex, and so too is $Z_- := Z - (\mathbb{R}_+ \times \mathbb{R}_+^N)$. Now, because (π^*, u^*) is Pareto optimal in Z, it is also Pareto optimal in Z_-, hence on the boundary of the latter. By the supporting hyperplane theorem, some nonzero $(\gamma, \lambda) \in \mathbb{R} \times \mathbb{R}^N$ exists such that

$$(\pi^*, u^*) \in \operatorname{argmax}_{(\pi,u)\in Z_-} [\gamma\pi + \lambda \cdot u] .$$

Because Z_- is downward comprehensive, the separation property implies $(\gamma, \lambda) \geqslant 0$. Scaling the nonzero nonnegative vector (γ, λ) by a strictly positive constant, we may assume without loss that $\max\{\gamma,\ \max_{i\in N} \lambda_i\} = 1$. Finally, the definition of Z implies $(\pi^* - 1, u^* + \vec{1}) \in Z$ too, so that $(\pi^*, u^*) \in \operatorname{argmax}_{(\pi,u)\in Z} [\gamma\pi + \lambda \cdot u]$ requires $\gamma \geqslant \vec{1} \cdot \lambda$. Thus, $\gamma = 1$, and λ is as desired. □

We now prove the characterization theorem.

Proof of Theorem 3. Lemma 13 says the two numbered conditions on ω are equivalent, so we need only show x^* is Pareto optimal if and only if some $\lambda \in \Lambda$ and λ-optimal ω exist such that $x^*(\boldsymbol{\theta}) = x_{\lambda,\omega}(\boldsymbol{\theta})$ almost surely.

First, suppose $\lambda \in \Lambda$, the vector ω is λ-optimal, and $x^*(\boldsymbol{\theta}) = x_{\lambda,\omega}(\boldsymbol{\theta})$ almost surely. By Lemma 15, then, the vector (π^*, u^*) is Pareto optimal in Z, where $\pi^* := \pi(x^*)$ and $u^* := \mathbb{E}\left[x^*(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \vec{\pi}^*$ Lemma 1 then implies that some transfer rule m^* exists such that (x^*, m^*) is IC and $\Pi(x^*, m^*) = \pi^*$; that $U(x^*, m^*) = u^*$ and (x^*, m^*) is IR; and (given revenue equivalence) that every alternative IC and IR mechanism (x, m) has $(\Pi(x, m), U(x, m)) \in Z$. Hence, Pareto optimality of the mechanism (x^*, m^*) follows from Pareto optimality of (π^*, u^*) in Z.

Conversely, suppose (x^*, m^*) is a Pareto optimal mechanism for some transfer rule m^*. Letting $\pi^* := \Pi(x^*, m^*)$ and $u^* := U(x^*, m^*)$, Lemma 1 and revenue equivalence tell us $(\pi^*, u^*) \in Z$. Lemma 14 therefore delivers some $(\tilde{\pi}, \tilde{u}) \geqslant (\pi^*, u^*)$ that is Pareto optimal in Z. By definition of Z, some allocation rule $\tilde{x}$ exists such that $\tilde{u} = \mathbb{E}\left[\tilde{x}(\boldsymbol{\theta})(\vec{b} - \boldsymbol{\theta})\right] - \tilde{\pi}\vec{1}$. Given Lemma 15, we may assume without loss that $\tilde{x} = x_{\lambda,\omega}$ for some $\lambda \in \Lambda$ and λ-optimal, and $\tilde{\pi} \leqslant \pi(\tilde{x})$. Because $\tilde{x}$ is monotone (hence interim monotone) and shifting transfers by a constant preserves IC, Lemma 1 tells us some transfer rule $\tilde{m}$ exists such that $(\tilde{x}, \tilde{m})$ is IC and generates $\Pi(x_{\lambda,\omega}, \tilde{m}) = \tilde{\pi}$, and that $(\tilde{x}, \tilde{m})$ is IR because $\tilde{\pi} \leqslant \pi(\tilde{x})$ and $U(\tilde{x}, \tilde{m}) = \tilde{u}$. Hence, the IC and IR mechanism $(x_{\lambda,\omega}, \tilde{m})$ generates a payoff vector $(\tilde{\pi}, \tilde{u}) \geqslant (\pi^*, u^*)$. Because the mechanism (x^*, m^*) is Pareto optimal, it follows that $(\tilde{\pi}, \tilde{u}) = (\pi^*, u^*)$. Finally, the uniqueness statement in Lemma 13 implies $x^*(\boldsymbol{\theta}) = \tilde{x}(\boldsymbol{\theta})$ almost surely, delivering the theorem. □

D.3.2. Generalizing other results to Pareto-optimal mechanisms

Now, we establish that the main result of Section 5 applies more generally to the entire Pareto frontier, as does the main result reported in Section C.1.

Proposition 7 (Simple mechanisms Pareto dominated)**:** *If $b < \bar{\theta}_j$ for every $j \in N$, then no collective posted-price mechanism is Pareto optimal, and no DIC mechanism is Pareto optimal.*

Proof. First, Theorem 3 tells us any Pareto-optimal allocation rule x is essentially identical to the (λ, ω)-allocation rule for some λ and ω. Because $\underline{\theta}_i < b < \bar{\theta}_i < \varphi_i(\bar{\theta}_i)$ for every $i \in N$, it follows that $0 < \mathbb{E}[x(\boldsymbol{\theta})] < 1$. Now, observe that x generates interim allocation rules X_i that are continuous on $(\underline{\theta}_i, \bar{\theta}_i)$ for every $i \in N$ with $\lambda_i + \omega_i < 1$, and nonconstant on $(\underline{\theta}_i, \bar{\theta}_i)$ if the optimal weights (λ, ω) have $\lambda_i + \omega_i > 0$. Indeed, the proof is identical to the proof of Lemma 5, but with Theorem 3 playing the role of Theorem 1, and $\lambda_j + \omega_j$ playing the role of ω_j and $\lambda_j \boldsymbol{\theta}_j + \omega_j \boldsymbol{\varphi}_j$ playing the role of $\omega_j \boldsymbol{\varphi}_j$ for each $j \in N$. To see that some agent i has X_i being both non-constant and continuous on $(\underline{\theta}_i, \bar{\theta}_i)$, it suffices to show no agent i has $\lambda_i + \omega_i = 1$; assume otherwise for a contradiction. Note that $\lambda \cdot \theta + \omega \cdot \varphi(\theta) = \omega_i \varphi(\theta_i) + (1-\omega_i)\theta_i$ is a strictly increasing transformation of $\theta_i \in \Theta_i$ that lies between θ_i and $\varphi_i(\theta_i)$. Therefore, some cutoff $p \in [\varphi_i^{-1}(b), b]$ exists such that $\lambda \cdot \boldsymbol{\theta} + \omega \cdot \boldsymbol{\varphi} \leqslant b$ if and only if $\boldsymbol{\theta}_i \leqslant p$. Hence,

$$\mathbb{E}[\boldsymbol{\varphi}_i \mid \lambda \cdot \boldsymbol{\theta} + \omega \cdot \boldsymbol{\varphi} \leqslant b] = \mathbb{E}[\boldsymbol{\varphi}_i \mid \boldsymbol{\theta}_i \leqslant p] = p,$$

where the last equality holds because a posted price of p (with agent i alone choosing whether to buy) generates the allocation rule $x_{\lambda,\omega}$ with binding IR for agent i. Therefore,

$$\mathbb{E}[\boldsymbol{\varphi}_i \mid \lambda \cdot \boldsymbol{\theta} + \omega \cdot \boldsymbol{\varphi} \leqslant b] \leqslant b < \bar{\theta}_j = \mathbb{E}[\boldsymbol{\varphi}_j] = \mathbb{E}[\boldsymbol{\varphi}_j \mid \lambda \cdot \boldsymbol{\theta} + \omega \cdot \boldsymbol{\varphi} \leqslant b],$$

as desired.

Next, we observe that no agent i exists such that $X_i^x|_{(\underline{\theta}_i, \bar{\theta}_i)}$ is both continuous and nonconstant, if (x, m) is either an IC collective posted-price mechanism or a bang-bang DIC mechanism—which will deliver the proposition. Given that we have seen $0 < \mathbb{E}[x(\boldsymbol{\theta})] < 1$, and that Theorem 3 tells us all Pareto-optimal mechanisms are bang-bang, the result follows directly from Lemma 11 for the case of bang-bang DIC mechanisms. So let us focus on showing it for the case of an IC posted-price mechanism. Let (x, m) be an IC collective posted price mechanism with price $p \in \mathbb{R}$. Below, we show no agent i exists such that $X_i^x|_{(\underline{\theta}_i, \bar{\theta}_i)}$ is both continuous and nonconstant. By the previous paragraph, it will follow that x is not Pareto optimal. Consider any agent i. Every $\theta_i, \hat{\theta}_i \in \Theta_i$ have

$$M_i^m(\hat{\theta}_i) - \theta_i X_i^x(\hat{\theta}_i) = (p - \theta_i) X_i^x(\hat{\theta}_i),$$

and so IC implies $X_i^x(\theta_i) = \max_{\hat{\theta}_i \in \Theta_i} X_i^x(\hat{\theta}_i)$ for any $\theta_i \in [\underline{\theta}_i, p)$ and $X_i^x(\theta_i) = \min_{\hat{\theta}_i \in \Theta_i} X_i^x(\hat{\theta}_i)$ for any $\theta_i \in (p, \bar{\theta}_i]$. In particular, X_i^x is constant both on $[\underline{\theta}_i, p)$ and on $(p, \bar{\theta}_i]$. Therefore, X_i^x is either constant on $(\underline{\theta}_i, \bar{\theta}_i)$ or discontinuous at $p \in (\underline{\theta}_i, \bar{\theta}_i)$. □

D.4. Proof for Section C.4

Proof of Proposition 6. First, consider the uniform-pricing regime. Following the agents' choice of shares σ', the buyer's optimal mechanism is characterized by Theorem 1. An optimal mechanism (x^*, m^*) is independent of the shares and depends only on the agents' distributions of types.[34] Because agents have identical distributions, the weights in the optimal mechanism are all $\omega_i = \frac{1}{N}$ and the optimal mechanism is symmetric. Hence, some $u \in \mathbb{R}_+$ exists such that a seller with share σ'_i gets payoff $\tau_i + \sigma'_i u$ if the buyer chooses this optimal mechanism and all sellers participate and truthfully report. The following play thus describes an equilibrium:

- The first seller proposes shares split σ and zero upfront transfers.
- Any other seller accepts a proposal $(\tilde{\sigma}, \tilde{\tau})$ if and only if $\tilde{\tau}_i + \tilde{\sigma}_i u > \sigma_i u$.
- For any realized shares σ', the buyer proposes the mechanism (x^*, m^*).
- If the mechanism (x^*, m^*) is proposed, then all sellers participate and truthfully report their types.
- If a mechanism other than (x^*, m^*) is proposed, then all sellers decline to participate.

This equilibrium has $\sigma' = \sigma$ and yields the buyer her optimal value, as required. Thus, the uniform-pricing game is pre-market non-manipulable.

Now, consider the discriminatory-pricing regime. Given any realized shares σ', seller i's cost of parting with his land is $\sigma'_i \boldsymbol{\theta}_i$, and so (a straightforward computation shows) his virtual cost is $\sigma'_i \boldsymbol{\varphi}_i$. Following Proposition 4.3 in Güth and Hellwig (1986), the essentially unique buyer-optimal mechanism for realized shares σ' has allocation rule given by $x(\theta) = \mathbb{1}_{\sigma' \cdot \boldsymbol{\varphi} \leqslant b}$, and transfers set so that IR binds for each agent.[35] By Proposition 4.2 of the same paper, seller i's expected payoff (gross of τ_i) is then $\mathbb{E}\left[\mathbb{1}_{\sigma' \cdot \boldsymbol{\varphi} \leqslant b}\ \sigma'_i(\boldsymbol{\varphi}_i - \boldsymbol{\theta}_i)\right]$. Therefore, the sum of the sellers' payoffs is

$$U(\sigma') := \mathbb{E}\left[\mathbb{1}_{\sigma' \cdot \boldsymbol{\varphi} \leqslant b}\ \sigma' \cdot (\boldsymbol{\varphi} - \boldsymbol{\theta})\right].$$

if a buyer-optimal mechanism is played—that is, if the buyer proposes it and all sellers participate and truthfully report—following share choice σ'.

To complete the proof of the proposition, we show by example that some specification of the model has $U(\tilde{\sigma}) > U(\sigma)$ for some $\tilde{\sigma} \in \Sigma$. The proposition will then follow, because any equilibrium would involve a successful proposal away from shares σ—for otherwise, the first seller could propose shares $\tilde{\sigma}$ together with lump-sum transfers to make every seller better off.

Consider the case with two sellers, each of whom has $\boldsymbol{\theta}_i$ uniform on $[0, 1]$, and

[34] The proof of Theorem 1 establishes that the optimal allocation rule is essentially unique, and Lemma 1 then implies all optimal mechanisms have the same interim transfer rules. In particular, all mechanisms yield the same per-share payoffs to all agents.

[35] Güth and Hellwig (1986) characterize profit-maximizing mechanisms for a seller who provides a public good to a group of agents and is allowed to use agent-specific transfers. A straightforward relabelling turns their model into one with a buyer who buys a public good from a group of sellers, so their analysis gives a characterization of buyer-optimal mechanisms. Güth and Hellwig (1986) impose a stronger regularity assumption (equivalent to assuming $\varphi_i(\theta_i) - \theta_i$ is increasing in our setting), but their proof applies identically under our weaker regularity assumption that φ_i is strictly increasing. Finally, they do not state essential uniqueness, but their proof establishes it because the allocation rule that solves their relaxed program is essentially unique.

a benefit $b = 1$. Then, any $\sigma' \in \Sigma$ has

$$U(\sigma') = \mathbb{E}\left[\mathbb{1}_{\sigma' \cdot \boldsymbol{\varphi} \leqslant b}\ \sigma' \cdot (\boldsymbol{\varphi} - \boldsymbol{\theta})\right] = \mathbb{E}\left[\mathbb{1}_{\sigma' \cdot \boldsymbol{\varphi} \leqslant 1}\ \sigma' \cdot \boldsymbol{\theta}\right].$$

In particular, the uniform share vector $\tilde{\sigma} = (\frac{1}{2}, \frac{1}{2})$ has

$$\begin{aligned} U(\tilde{\sigma}) &= \tfrac{1}{2}\mathbb{E}\left[\mathbb{1}_{\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 \leqslant 1}\ (\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2)\right] \\ &= \mathbb{E}\left[\mathbb{1}_{\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 \leqslant 1}\ \boldsymbol{\theta}_1\right] \\ &= \int_0^1 \int_0^{1-\theta_1} \theta_1\ \mathrm{d}\theta_2\ \mathrm{d}\theta_1 = \int_0^1 (1 - \theta_1)\theta_1\ \mathrm{d}\theta_1 = \left[\tfrac{1}{2}\theta_1^2 - \tfrac{1}{3}\theta_1^3\right]_{\theta_1=0}^1 \\ &= \frac{1}{6}. \end{aligned}$$

Meanwhile, as $\sigma \to (1, 0)$, the quantity $U(\sigma)$ converges (by the dominated convergence theorem) to

$$\mathbb{E}\left[\mathbb{1}_{\boldsymbol{\theta}_1 \leqslant \frac{1}{2}}\ \boldsymbol{\theta}_1\right] = \int_0^{\frac{1}{2}} \theta_1\ \mathrm{d}\theta_1 = \left[\tfrac{1}{2}\theta_1^2\right]_{\theta_1=0}^{\frac{1}{2}} = \frac{1}{8} < \frac{1}{6}.$$

In particular, when the initial shares σ are sufficiently asymmetric, we have $U(\tilde{\sigma}) > U(\sigma)$, as required. □