

SELLING TO A GROUP

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

July 11, 2021

ABSTRACT

A group of agents can collectively purchase a public good that yields heterogeneous benefits to its members. Combining a reduced-form implementation result with a duality argument, we characterize the seller’s profit-maximizing mechanism. Trade outcomes depend solely on a weighted average of the agents’ virtual values, with endogenous voting weights. Heterogeneity in voting weights reflects heterogeneity in agents’ value distributions, where agents with lower value distributions are given more weight in trade decisions. Simple pricing rules are generally not (even approximately) optimal.

1. Introduction

We study a problem of selling a good to a *buyer group*, consisting of multiple agents. The problem has two key features. First, the good is public, that is, conditional on sale, its benefits are enjoyed by all the group members. Second, the purchase is financed from a collective pool of money at the group’s disposal. Examples with these features abound: a software company that needs to convince a committee of senior managers in an organization to purchase its product, a consultant who sells her proposals to the members of an executive board, or a contractor who requires approval from a city council’s members to fund a public project.

Our main goal is to characterize the seller’s optimal (profit-maximizing) mechanisms when group members’ benefits from the good are private information, and each group member can individually veto the mechanism. The example below demonstrates the key forces in our environment and the qualitative insights that emerge from our characterization.

*haghpanah@psu.edu

†a.kuvalekar@essex.ac.uk

‡e.lipnowski@columbia.edu

§We thank Nageeb Ali, Ben Brooks, Rahul Deb, Navin Kartik, Benny Moldovanu, and Doron Ravid for helpful discussions.

EXAMPLE 1: Two agents, 1 and 2, have independently drawn values for the good. Agent 1’s value, θ_1 , is distributed uniformly over $[0, 2]$, whereas agent 2’s value, θ_2 , is distributed uniformly over $[0, 3]$. The seller can provide at most one unit of the good at zero cost. The seller designs a mechanism—an allocation rule and a transfer rule—that conditions on agents’ reports of their respective types. If a sale occurs and the common transfer is m , each agent i ’s utility is $\theta_i - m$; if no sale occurs and the transfer is m , each agent’s utility is $-m$. Agents’ outside option from not participating (i.e., vetoing the mechanism) is zero. ■

Let us begin with the simple class of posted-price mechanisms. Such mechanisms would be optimal, following classic results, if the seller interacted with a single agent. Because agent 2’s valuation first-order stochastically dominates that of agent 1, a reasonable starting point is the optimal posted price for agent 2. That is, the seller could post a price of $\frac{3}{2}$ and make a take-it-or-leave-it offer to agent 2. The sale takes place if agent 2 agrees, but not otherwise. However, this mechanism yields a negative (interim) expected payoff to agent 1 when his type is low, and so this agent will want to veto the mechanism. For illustration, note that if, for example, $\theta_1 = 0$, the good will be sold at a price of $\frac{3}{2}$ with a positive probability, thereby earning agent 1 a negative interim payoff. Modifying the payment rule to make this posted-price mechanism individually rational (IR) is not difficult. The seller can offer a subsidy: an additional transfer that she pays the agents regardless of whether a sale occurs. The subsidy must be at least as high as the expected revenue itself to make the mechanism IR, in particular, when $\theta_1 = 0$. Therefore, the seller earns at most zero profit from such a mechanism.

One remedy to the above problem is to require *unanimous* approval of the trade. That is, the good is offered at a price p , and is sold if and only if both agents agree to the purchase. We call this rule a “unanimous posted price” mechanism. This mechanism is IR by construction, and generates a profit of $p(1 - \frac{p}{2})(1 - \frac{p}{3})$ —maximized at $p = \frac{1}{3}(5 - \sqrt{7})$. This maximal profit is approximately 0.35.

Can the seller do better? Her problem is to maximize profit over all (Bayesian) incentive compatible (IC) and IR mechanisms. If she could use agent-specific transfers, standard arguments à la Myerson (1981) teach us a given allocation rule is implementable if and only if its associated interim allocations are nondecreasing. Perhaps surprisingly, even with access only to collective transfers, Lemma 2 shows the same condition characterizes implementability in our setting. However, absent agent-specific transfers, the agents’ average payments must coincide. Hence, the maximal revenue that can be extracted from a given implementable allocation rule is pinned down by the condition that *one* agent’s individual rationality constraint binds (and the others’ are satisfied). Therefore, the profit-maximizing mechanism given an allocation rule $x(\cdot)$ yields a profit of $\min_{i \in N} \mathbb{E}\{x(\boldsymbol{\theta})[\varphi_i(\boldsymbol{\theta}_i) - c]\}$, where N is the set of agents, $\varphi_i(\cdot)$ is the virtual valuation function of agent i , and c is the production cost of the good to the seller.

If we temporarily ignore the above monotonicity condition on allocation rules,

we can view the seller’s optimal profit,

$$\max_x \min_{i \in N} \mathbb{E}\{x(\boldsymbol{\theta})[\varphi_i(\boldsymbol{\theta}_i) - c]\},$$

as the maximin value from a two-player zero-sum game in which one player, the Maximizer, chooses an allocation rule $x(\cdot)$, and the other player, the Minimizer, chooses an agent i . We show this infinite game has a Nash equilibrium, and hence its value is

$$\min_{\omega \in \Delta N} \max_x \mathbb{E}\{x(\boldsymbol{\theta})[\omega \cdot \varphi(\boldsymbol{\theta}) - c]\},$$

where ω , a distribution over agents, can also be seen as a mixed strategy for the Minimizer. We characterize the equilibria of the above game to establish that the optimal allocation rule is unique and is a weighted voting rule: the good is sold if and only if $\omega \cdot \varphi(\boldsymbol{\theta}) \geq c$, where the optimal weights ω are characterized by the simple program, $\min_{\omega \in \Delta N} \mathbb{E}\{[\omega \cdot \varphi(\boldsymbol{\theta}) - c]_+\}$. Moreover, under a regularity assumption, this allocation rule is monotone and hence solves the seller’s problem. These results are summarized in Theorem 1.

Solving the above problem for Example 1, we obtain that the optimal weights are $\omega_1 = \sqrt{\frac{3}{7}}$ and $\omega_2 = 1 - \omega_1$ and the associated optimal profit is approximately 0.39. Notice how the allocation rule of the optimal mechanism, where the good is sold if and only if $\omega_1 \boldsymbol{\theta}_1 + \omega_2 \boldsymbol{\theta}_2 \geq 1 + \frac{\omega_2^2}{2}$, is different from the allocation rule of the unanimous posted-price mechanism, where the good is sold if and only if $\min(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \geq p$. The optimal allocation rule, with its interior ω , relaxes the stringent requirement of a unanimous agreement.

Consistent with Example 1, generally, the optimal mechanism overweights agents with lower value distributions. In fact, Theorem 2 delivers a quantitative ranking result for the optimal voting weights whenever the optimal mechanism is nontrivial. It says that, if φ_i is smaller than $\alpha \varphi_j$ in the hazard-rate order (a strong form of stochastic dominance that has been previously applied in the auction literature) for some $\alpha \in (0, 1]$, the optimal mechanism entails $\omega_i \geq \alpha \omega_j$. The example exhibits this distributional ranking with $\alpha = \frac{3}{2}$; and indeed, a direct computation shows $\omega_1 > \frac{3}{2} \omega_2$. Qualitatively, the seller listens more to weaker agents, whose veto constraint binds more tightly. The same lesson is reflected in Corollary 2, which shows the optimal mechanism uses only one agent’s private information if and only if that agent’s value distribution is extremely low relative to others’.

Though unanimous posted prices are not optimal (as we saw above for the example), we do find they are profit-maximizing within a natural classes of simple mechanisms—specifically mechanisms for which the price of the good is constant conditional on trade (Proposition 2). As formalized in Proposition 3, posted-price mechanisms are suboptimal in all interesting instances of our model. We further quantify this payoff loss in Proposition 5, by considering a many-agent limit case of our model. There, we show the “price of simplicity”—the limiting ratio of the

optimal profit to that of unanimous posted pricing—is unboundedly large except in asymptotically trivial cases.

Finally, we turn to exploring alternative bargaining arrangements in Section 6. Although we focus on seller-optimal mechanisms throughout, the proof approach for Theorem 1 can be adapted to fully characterize the Pareto frontier. In particular, as Theorem 3 shows, any Pareto-optimal mechanism allocates the good if and only if a weighted average of virtual values and true values exceeds the cost.

We also explore the consequences of relaxing another important feature of our environment, namely, that any agent can unilaterally veto the mechanism. Whereas smaller groups often do assign veto rights to their members, a more permissive bargaining arrangement may be better suited for understanding some contexts. For example, rather than unanimity, the group might require that the mechanism be approved by some minimum number of agents. Although intuitively appealing, capturing such flexible arrangements in a reasonable framework seems elusive. We discuss the associated issues in some detail in Section 6.2, but the broad takeaway is that we are not aware of an appropriate model that relaxes veto bargaining while retaining the spirit and tractability of our main analysis. We see the pursuit of such a model as an exciting avenue for future research.

1.1. Related Work

Because the good for sale in our model is public, our work is closely related to the vast 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. Several papers show, in related contexts, that any mechanism achieving ex-ante budget balance can be converted (preserving agents’ incentives) to a mechanism with ex-post budget balance by choosing ex-post transfers appropriately (e.g., Makowski and Mezzetti, 1994; d’Aspremont et al., 2004; Börgers and Norman, 2009). Our construction of ex-post transfer rules that induce a given profile of interim transfer rules is related, especially in the two-agent special case. Rob (1989) shows that with a large number of agents, profit-maximizing mechanisms are inefficient, whereas Mailath and Postlewaite (1990) extend this inefficiency result to all IR and budget-balanced mechanisms. In a setting where agents’ values for a good are symmetric, and each is initially endowed with a share, Cramton et al. (1987) show efficient and IR trading mechanisms exist if and only if agents’ shares are sufficiently symmetric. Güth and Hellwig (1986) identify profit-maximizing mechanisms subject to incentive compatibility and individual rationality constraints. Hence, our seller’s problem is equivalent to that of Güth and Hellwig (1986), with the added restriction that agent-specific transfers are not available.

Another strand of the literature on public goods studies voting mechanisms without monetary transfers. Starting with Rae (1969), many entries to 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 IC mechanism must be a weighted-majority rule, and characterize the weights that maximize efficiency.¹ Our model is a middle ground between the two aforementioned strands of literature on public goods, in that monetary transfers are available in our setting but are restricted to be identical across agents.

Our work is also related to the literature that studies (approximate) optimality of posted-price mechanisms. Myerson (1981) and Riley and Zeckhauser (1983) show posted pricing is an optimal strategy for selling a single good to a single agent. Even though posted pricing is no longer optimal in settings with multiple goods or agents, it remains approximately optimal in many such settings (see, e.g., Chawla et al., 2010; Chawla et al., 2015; Hart and Nisan, 2017; Babaioff et al., 2020). By contrast, in our setting, posted-price mechanisms perform arbitrarily poorly relative to optimal mechanisms as the number of agents grows.

We solve for optimal mechanisms using a simple reduced-form characterization of implementable collective transfer rules. Our work is thus related to the literature on reduced-form implementation in auctions (e.g., Border, 1991; Cai et al., 2012; Che et al., 2013; Alaei et al., 2019). Our implementability result could be repurposed to study interim allocation rules for a real-valued (or nonnegative real-valued) and unbounded public outcome. This explicit, tractable implementability result for transfer rules stands in contrast to the results of Gopalan et al. (2018), who show that if the public outcome is binary-valued (or, equivalently under linear preferences, if it is restricted to some bounded interval), no computationally tractable characterization of implementable reduced forms of collective transfer rules exists.

Our work is related to the literature on the (in)equivalence of Bayesian and dominant strategy incentive compatibility. Incentive compatibility in our setting with a single good is characterized by standard monotonicity constraints (as in Myerson, 1981). Nonetheless, because individual transfers are not permitted, optimal mechanisms are not dominant-strategy incentive compatible in our setting, except in uninteresting cases. This result stands in contrast to the known results on the equivalence of Bayesian and dominant-strategy incentive compatibility in settings with unidimensional types and agent-specific transfers (Manelli and Vincent, 2010; Gershkov et al., 2013).²

¹Also see Gershkov et al. (2017), who further study optimal voting mechanisms for a class of environments with more than two social outcomes.

²Such equivalence is known to fail in the case of multidimensional private information (Jehiel et al., 1999; Gershkov et al., 2013; Yao, 2017; Manelli and Vincent, 2019).

2. Model

We study the problem of a seller who can sell one indivisible good to be shared by a group of agents. We denote the finite nonempty set of agents $N = \{1, \dots, N\}$. The seller incurs a cost $c \geq 0$ if the good is sold. Any monetary transfer paid for the good is borne collectively by the group. Agents are heterogeneous in how they value the good vis-à-vis the group's money. That is, each agent i has a private type θ_i , which is a random variable taking values in $\Theta_i = [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}_+$. Agents' types are independent, and i 's type follows the cumulative distribution function F_i .³ We make the following regularity assumption for each $i \in N$: the CDF F_i admits a continuous and strictly positive density f_i on its support, and the *virtual value* $\varphi_i : \Theta_i \rightarrow \mathbb{R}$ given by $\varphi_i(\theta_i) := \theta_i - \frac{1-F_i(\theta_i)}{f_i(\theta_i)}$ is strictly increasing. Working directly with an agent's virtual value $\varphi_i := \varphi_i(\theta_i)$, an atomlessly distributed random variable with convex support, will often be convenient.

An outcome of our contracting environment consists of a probability $x \in [0, 1]$ with which the good is sold to the buyer group, and a (signed) transfer $m \in \mathbb{R}$ paid to the seller by the buyer group. The payoff of agent i for this outcome is given by $x\theta_i - m$, whereas the seller's payoff is $m - cx$. Let us highlight two distinguishing features of our environment. First, the good is public: conditional on it being allocated, every agent derives a benefit (equal to his type). Second, the transfers are collective. One could interpret our agents as a group of committee members deciding whether to approve a project (purchased from our seller). The transfer paid for the project will come from the common pool of money that the committee can access, whereas the private benefits that each member derives from using the organization's funds on this particular project may vary.

The seller knows the distribution of types for each agent but not agents' realized types. She designs a (direct, without loss) mechanism, which specifies a probability of trade and a total transfer, for every profile of reported types. For most of the paper, we focus on the optimal (i.e., profit-maximizing) mechanism for the seller, among all mechanisms that are IC and IR for the agents.⁴ We formally define these standard notions below.

DEFINITION 1: A *(collective) allocation rule* is a measurable function $x : \Theta \rightarrow [0, 1]$; let \mathcal{X} denote the set of all allocation rules. A *(collective) transfer rule* is a bounded measurable function $m : \Theta \rightarrow \mathbb{R}$. A *(collective selling) mechanism* is a pair (x, m) consisting of an allocation rule and a transfer rule.

³We 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 CDF interchangeably, and use F and F_{-i} to refer to associated product measures on Θ and Θ_{-i} , respectively.

⁴In Section 6, we consider alternative bargaining arrangements. For instance, we consider the range of Pareto-optimal mechanisms (not only seller-optimal ones) to account for settings in which the seller has only imperfect bargaining power, and we consider agent approval processes beyond veto bargaining, discussing how one might relax the below individual rationality constraint.

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

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

and say it is **individually rational (IR)** if

$$\mathbb{E} [\theta_i x(\theta_i, \boldsymbol{\theta}_{-i}) - m(\theta_i, \boldsymbol{\theta}_{-i})] \geq 0, \quad \forall i \in N, \quad \forall \theta_i \in \Theta_i. \quad (\text{IR})$$

The **profit** generated by a mechanism (x, m) is $\Pi(x, m) := \mathbb{E} [m(\boldsymbol{\theta}) - cx(\boldsymbol{\theta})]$. An **optimal mechanism** is an IC and IR mechanism that generates a weakly higher profit than any other IC and IR mechanism. Finally, an **optimal allocation rule** is any allocation rule x such that (x, m) is an optimal mechanism for some m .

3. Characterizing the Optimal Mechanism

In this section, we fully characterize optimal mechanisms. First, we provide a useful reduced-form implementation result for transfers, characterizing exactly which profiles of interim transfer rules can be implemented with some collective transfer rule. Then, using this characterization, we describe which allocation rules are implementable, and solve for the seller's optimal profit from implementing such an allocation rule. Next, we establish that a unique optimal allocation rule exists and can be described as a weighted voting rule with weights that we explicitly characterize. Finally, we show that, except in trivial cases, the voting weights that describe an optimal allocation rule are unique; hence, characterizing them is equivalent to characterizing optimal mechanisms.

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 2: Fix any agent $i \in N$. Given an allocation rule x , define the **interim allocation rule** to be $X_i^x : \Theta_i \rightarrow \mathbb{R}$ 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 $M_i^m : \Theta_i \rightarrow [0, 1]$ 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 increasing for every $i \in N$.

As a first step toward solving our seller's problem, we provide a simple reduced-form implementation result, which exactly characterizes which profiles of interim transfer rules can be induced when transfers are restricted to be ex-post identical for all agents. The average transfer stipulated by each interim transfer rule must be the same, both being the expected value (by iterated expectations) of a common random variable. The following lemma, which may be of independent interest,

shows this necessary condition is sufficient. Moreover, it shows that if transfers are restricted to be nonnegative, pairing that necessary condition with the another obviously necessary condition (that all interim transfers be nonnegative) is again sufficient for implementability.

LEMMA 1 (Reduced-form transfer rules): *Let $M_i^* : \Theta_i \rightarrow \mathbb{R}$ be a bounded measurable function for each $i \in N$. Then, the following are equivalent:*

1. *Some transfer rule m exists such that each $i \in N$ has $M_i^m = M_i^*$.*
2. *The expectations $\{\mathbb{E}[M_i^*(\theta_i)]\}_{i \in N}$ all coincide.*

Moreover, m can be taken to be nonnegative if and only if each of $\{M_i^\}_{i \in N}$ is.*

The straightforward proof of the above lemma is constructive and resembles previous constructions in the literature (e.g., Makowski and Mezzetti, 1994; d’Aspremont et al., 2004; Börgers and Norman, 2009) that convert ex-ante budget-balanced mechanisms into ex-post budget-balanced mechanisms, while preserving the players’ interim transfer rules. Although we apply this result to collective transfers, we imagine future applications could benefit from Lemma 1, treating its nonnegative version as a reduced-form implementability result for (unbounded) public good provision.

Leveraging the above result, the following lemma characterizes allocation rules that are implementable with some transfer rule, as well as the seller’s profit from implementing such an allocation rule.

LEMMA 2 (Implementable allocations): *Let x be some allocation rule.*

1. *Mechanism (x, m) is IC and IR for some transfer rule m if and only if x is interim monotone.*
2. *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*

$$\min_{i \in N} \mathbb{E}[x(\theta)(\varphi_i - c)].$$

Classic results (Myerson, 1981) would imply interim monotonicity fully characterizes implementability, if the seller could freely choose the interim transfer rule that each agent faces. However, our seller is constrained in that different agents’ interim transfers must be derived from a common ex-post transfer rule. Nevertheless, Lemma 1 tells us the sole constraint that collective transfers place on these interim transfer rules is that they stipulate the same transfer *on average*. Hence, after modifying the transfer rules by a player-dependent flat subsidy (which does not affect IC), they can be implemented by some ex-post transfer rule. Consequently, interim monotonicity fully characterizes implementability of an allocation rule, as in the setting with separable payments.

Given that an allocation rule is implementable, the reasoning behind Myerson’s (1981) result determines each agent’s interim transfer rule up to a constant.

However, that each agent must pay the same transfer on average (because they do so ex post) determines the entire profile of such constants up to a single scalar parameter. Analogous to how an optimal auction would optimize the transfer rule by setting each agent’s IR constraint to bind, our remaining constant is solved out by imposing that *one* agent’s IR binds (and the others’ are satisfied). Hence, in contrast to the implementability question, the seller’s maximum (IC and IR) profit from a given allocation rule is affected by the fact that interim transfer rules cannot be separably designed.

With Lemma 2 in hand, our seller’s problem can be recast directly as an optimization over allocation rules, with the associated profit of such a rule being pinned down by revenue equivalence and the principle that IR binds for the worst-off low type. Formally, the seller’s optimization over allocation rules is

$$\begin{aligned} \max_{x \in \mathcal{X}} \left\{ \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\varphi_i - c)] \right\} & \quad (\text{SP}) \\ \text{s.t. } x \text{ is interim monotone.} & \end{aligned}$$

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

DEFINITION 3: *Given $\omega \in \Delta N$, the ω -voting rule is the allocation rule $x_\omega := \mathbb{1}_{\omega \cdot \boldsymbol{\varphi} \geq c}$. Say an allocation rule is a **voting rule** if it is a ω -voting rule for some $\omega \in \Delta N$.*

We now state our main characterization theorem.

THEOREM 1 (Optimal allocation): *An essentially unique optimal allocation rule exists and is a voting rule.⁵ Namely, the ω -voting rule is optimal for any $\omega \in \Delta N$ that satisfies either of the following two equivalent conditions (and some such ω exists):*

1. $\omega \in \operatorname{argmin}_{\tilde{\omega} \in \Delta N} \mathbb{E}[(\tilde{\omega} \cdot \boldsymbol{\varphi} - c)_+]$.
2. $\operatorname{supp}(\omega) \subseteq \operatorname{argmin}_{i \in N} \mathbb{E}[\varphi_i \mathbb{1}_{\omega \cdot \boldsymbol{\varphi} \geq c}]$.

Moreover, an optimal mechanism exists with nonnegative transfers.

The proof of Theorem 1 studies a relaxed program (RSP) 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 IC and IR constraints must be satisfied, and the objective of the game is $\mathbb{E}[x(\boldsymbol{\theta})(\varphi_i - c)]$ —the seller’s highest possible profit from the chosen allocation, subject to the “revenue equivalence” formula and the chosen agent’s IR constraint.⁶ Observe that an allocation rule

⁵By “essentially unique,” we mean any alternative optimal allocation rule x has $x(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely.

⁶So if the Maximizer’s chosen allocation rule happens to be interim-monotone, the objective

solves (RSP) if and only if it is a cautious optimum for the Maximizer in the auxiliary game, that is, a “maximin” strategy. Moreover, standard results on 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 a minimax theorem (for infinite games), and every mixed strategy ω for the Minimizer exhibits a unique (up to almost-everywhere equality) best response x_ω for the Maximizer. 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. Specifically, this allocation rule is equal to the ω -weighted sum of agents’ virtual values, where $\omega \in \Delta N$ is a Nash equilibrium strategy for the Minimizer. The pair of conditions characterizing such ω are standard to zero-sum games: the mixed strategy ω is a cautious optimum for the Minimizer (the first condition) if and only if it is a best response to some Maximizer best response to ω (the second condition, once 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 values are increasing, a cutoff rule for the ω -weighted virtual value is monotone, hence interim monotone. The result is a characterization of the unique optimal allocation rule, solving not only (RSP) but also (SP).

Finally, we turn to the form of optimal transfer rules. Having solved for the optimal allocation rule and the expected revenue that the seller garners, each agent’s *interim* expected transfer rule is fully determined by the classic Myerson (1981) payment formula. Moreover, direct computation shows these interim transfers are always nonnegative; that is, no agent expects (even conditioning on realizing his lowest possible type) to be subsidized on average. Although infinitely many ex-post transfer rules implement these interim transfers, and some will indeed specify a negative payment for some type profile realizations, Lemma 1 shows by construction that at least one such transfer rule does not. The theorem follows.

Given Theorem 1, we can characterize optimal mechanisms by characterizing which voting weights solve the two equivalent conditions listed in the theorem. This goal justifies the following definition.

DEFINITION 4: Say $\omega \in \Delta N$, is an *optimal* vector of voting weights if it is in $\operatorname{argmin}_{\tilde{\omega} \in \Delta N} \mathbb{E}[(\tilde{\omega} \cdot \varphi - c)_+]$ or, equivalently, has $\mathbb{E}[\varphi_i \mathbb{1}_{\omega \cdot \varphi \geq c}] \leq \mathbb{E}[\varphi_j \mathbb{1}_{\omega \cdot \varphi \geq c}]$ for every $i, j \in N$ with $\omega_i > 0$.

In light of Theorem 1 (together with revenue equivalence), understanding optimal selling mechanisms amounts to understanding which voting weights $\omega \in \Delta N$

is the seller’s highest possible profit from the chosen allocation, subject to IC and IR for the Minimizer’s chosen agent.

are optimal.⁷

Whereas Theorem 1 delivered the uniqueness of the optimal allocation rule, it is straightforward to see that an optimal ω need not be unique in trivial cases. For example, if $\max_{i \in N} \bar{\theta}_i \leq c$, observe that every choice $\omega \in \Delta N$ of voting weights is optimal, each inducing a mechanism in which trade never occurs. Similarly, optimal weights may not be unique when the optimal mechanism stipulates that trade always occurs. This observation motivates the following definition.

DEFINITION 5: The *never-trade* mechanism is given by $(x, m) = (0, 0)$. The *always-trade* mechanism is given by $(x, m) = (1, \min_{i \in N} \underline{\theta}_i)$. A mechanism (x, m) or an allocation rule x is **trivial** if $\mathbb{E}[x(\boldsymbol{\theta})] \in \{0, 1\}$ and **nontrivial** if $0 < \mathbb{E}[x(\boldsymbol{\theta})] < 1$.

The next result shows the above multiplicity happens only when either never trading or always trading is optimal, and characterizes when these cases arise. In all other (and so in all interesting) cases, the optimal ω is unique.

PROPOSITION 1 (Trivial optimal mechanisms and uniqueness of weights):

1. The never-trade mechanism is optimal if and only if $\min_{j \in N} \bar{\theta}_j \leq c$.
2. The always-trade mechanism is optimal if and only if some $i \in \operatorname{argmin}_{j \in N} \underline{\theta}_j$ exists such that $\varphi_i(\underline{\theta}_i) \geq c$.
3. In all other cases, a nontrivial mechanism is optimal and a unique $\omega \in \Delta N$ exists such that x_ω is optimal.

Because the never-trade and always-trade mechanisms are obviously optimal among mechanisms that never or always allocate, respectively, the real content of the first two statements of the proposition is a characterization of when each of the $x = 0$ and $x = 1$ allocation rules is optimal.⁸ Given Theorem 1, it suffices to check when an optimal voting-weight vector exists that generates each of these allocations—a straightforward computation.

The third statement provides further comfort in directly interpreting ω as voting weights, because it delivers (in all nontrivial cases) a one-to-one correspondence between the optimal allocation rule and the weights that describe it. Using that the profile of virtual values has convex support, the proof shows that distinct voting weights that generate nontrivial allocation rules must generate distinct allocations with positive probability. Hence, the result follows from the uniqueness part of Theorem 1.

⁷As a trivial observation, note an optimal (indirect) mechanism exists in which each bidder submits a vote from a bounded interval, with the good being provided if and only if the weighted sum of votes exceeds a threshold.

⁸Although Theorem 1 enables a unified simple proof of this and related results, note the characterization of when the never-trade mechanism is optimal is easy to establish directly. If $\bar{\theta}_i \leq c$ for some $i \in N$, one can show agent i 's IR constraint implies $\mathbb{E}[m(\boldsymbol{\theta})] \leq c \mathbb{E}[x(\boldsymbol{\theta})]$, precluding positive profit. If $\min_{i \in N} \bar{\theta}_i > c$, posting a price strictly between these two quantities, and requiring unanimous agreement to buy, generates positive profit.

4. The Role of Heterogeneity

Because optimal mechanisms take the form of a weighted voting rule, considering the relative voting weights of different agents is natural. This section asks which agents are assigned a high weight in determining the trade decision.

In a relaxed program in which the seller can observe agents' types (but is still subject to IR), trade occurs if and only if the lowest-value agent's type exceeds the production cost. Hence, a reasonable intuition is that agents with the lowest value distribution (in some sense) will be overweighted in the decision. Indeed, examining whether the extreme always-trade and never-trade mechanisms are optimal shows that when one agent has a stochastically lowest value distribution, that agent alone determines the (sub)optimality of said mechanism.

COROLLARY 1 (Low-value agents and the extensive margin): *Suppose $i \in N$ is such that θ_i is (weakly) first-order stochastically dominated by θ_j (i.e., $F_i \geq F_j$) for each $j \in N$. Then, the never-trade [resp. always-trade] mechanism is optimal if and only if $\bar{\theta}_i \leq c$ [resp. $\varphi_i(\underline{\theta}_i) \geq c$].*

The above corollary follows directly from Proposition 1, once one observes that an agent with a (first-order stochastically) lowest value distribution necessarily has a lowest high type, a lowest low type, and a lowest low virtual valuation among those who have the lowest low type.

Although Corollary 1 provides a sense in which a weak agent is pivotal to the nature of the allocation, it is a weak result. In particular, conditional on one of these trivial mechanisms being used, all agents are treated equally—facing the same ex-post outcome and all having their type realization ignored. Understanding, more generally, when an agent is pivotal to determining the allocation would be desirable. A particularly strong notion of pivotality is captured by the following definition.

DEFINITION 6: *Given an agent $i \in N$, let i -dictatorship refer to the vector $\omega \in \Delta N$ of voting weights in which all agents other than i are ignored, that is, with $\omega_j = \mathbb{1}_{i=j}$ for each $j \in N$.*

The following corollary characterizes when an i -dictatorship mechanism is optimal for the seller. Such an allocation rule can be implemented by posting a price $p \in [\underline{\theta}_i, \bar{\theta}_i]$ that ensures trade occurs if and only if i 's virtual value exceeds the production cost. The next result, a nearly immediate consequence of Theorem 1, shows such a mechanism is optimal if and only if this price is such that every type of every other agent would happily trade at that price.

COROLLARY 2 (Dictatorship): *Given $i \in N$, the following are equivalent:*

1. *The i -dictatorship voting rule is optimal.*
2. *Either no price $p \in (c, \bar{\theta}_i)$ exists, or the optimal posted price p when facing only i satisfies $p \leq \underline{\theta}_j$ for every agent $j \neq i$.*

3. Either $\bar{\theta}_i \leq c$, or $\bar{\theta}_i > c$ and $\mathbb{E}[\varphi_j] \geq \mathbb{E}[\varphi_i | \varphi_i \geq c]$ for every agent $j \neq i$.

Hence, the case of dictatorship is very special, requiring extreme asymmetry between the agents. If the optimal mechanism is nontrivial, observe i -dictatorship violates IR if other agents have the same lowest type as agent i (or lower). The next result strengthens this observation, suggesting the typical optimal mechanism pays some attention to all agents. Indeed, in the nontrivial case, if the lowest possible value is zero for each agent, the result implies the unique optimal allocation is responsive to every agent's private information.

COROLLARY 3 (Not ignoring the lowest types): *Suppose the optimal allocation rule is nontrivial. If $i \in N$ has $\underline{\theta}_i \leq \underline{\theta}_j$ for every $j \in N$, then the unique optimal voting weights $\omega \in \Delta N$ have $\omega_i > 0$. In particular, if $\underline{\theta} = \vec{0}$, then $\omega_i > 0$ for every $i \in N$.*

The proof of the above corollary is nearly immediate from Theorem 1. In the auxiliary zero-sum game that characterizes optimal allocation rules, we show that any agent whose lowest type is (weakly) lower than everybody else's will necessarily be a unique best response for the Minimizer to any allocation rule that ignores his type.

The previous results of this section have all spoken to the choice of which agents will exert *some* influence over the eventual trade decision in the optimal mechanism, but they have been silent on the degree of such influence. For the remainder of this section, we pursue a quantitative analysis of the optimal voting weights. Specifically, we seek conditions on primitives under which we can rank ω_i and ω_j for two agents i and j (and under which we can quantify a wedge between these two weights). To state our main condition, we invest in the following distributional ranking definition.

DEFINITION 7: *Given two real random variables \mathbf{v} and \mathbf{w} with respective CDFs given by G and H , say \mathbf{v} is smaller than \mathbf{w} in the hazard-rate order if $\sup[\text{supp}(\mathbf{v})] \leq \sup[\text{supp}(\mathbf{w})]$, and $\frac{1-H}{1-G}$ is weakly increasing on $(-\infty, \sup[\text{supp}(\mathbf{v})])$.*⁹

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 above 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 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.¹⁰ In our set-

⁹If \mathbf{v} and \mathbf{w} both admit continuous densities on their supports, the monotonicity of the function $\frac{1-H}{1-G}$ can be equivalently expressed (taking a derivative) as requiring that the hazard rate of \mathbf{v} is weakly below that of \mathbf{w} on the same interval.

¹⁰These papers rank bidders' value distributions according to the *reverse* hazard-rate order, which, by Theorem 1.B.41 of Shaked and Shanthikumar (2007), amounts to a hazard-rate or-

ting, as the following theorem shows, a hazard-rate order on agents' *virtual* value distributions is of relevance in designing optimal selling mechanisms.

THEOREM 2 (Ranking voting weights): *If φ_i is smaller than $\alpha\varphi_j$ in the hazard-rate order, where $\alpha \in (0, 1]$, and the optimal allocation rule is nontrivial, then the unique optimal vector of voting weights ω satisfies $\omega_i \geq \frac{1}{\alpha}\omega_j$.¹¹*

We prove the theorem by contradiction, assuming some optimal voting weights fail to satisfy the desired ranking property. Modifying the weights on agents i and j , we construct an alternative vector of voting weights for the seller to use. Importantly, the original weight vector is more assortative with the virtual value distributions than the rearranged weight vector. Hence, we can apply known results that translate hazard-rate orders on random variables to increasing convex orders on their weighted sums, inferring that the rearranged weight vector must also be optimal. Although the new optimum we construct may not itself satisfy the α -ranking property that the theorem requires, we still derive a contradiction with the uniqueness of optimal voting weights that Proposition 1 guarantees. The theorem follows.

We conclude the section with two different classes of examples to which Theorem 2 applies. One is a rigid ranking of value distributions, where the distribution of a given agent's values can be arbitrary (subject to our standing regularity conditions) but different agents' distributions are assumed proportional. The other is the full class of uniform distributions. In either case, we note the agents' distributions of virtual values inherits the same structure, making the proposition straightforward to apply.

EXAMPLE 2 (Proportional value distributions): *Suppose $\theta_i/\bar{\theta}_i$ and $\theta_j/\bar{\theta}_j$ are identically distributed; for example, this property would hold if θ_i and θ_j were uniformly distributed on $[0, \bar{\theta}_i]$ and $[0, \bar{\theta}_j]$, respectively. A direct computation shows $\varphi_i/\bar{\theta}_i$ and $\varphi_j/\bar{\theta}_j$ are identically distributed too in this case. Theorem 2 then implies, if $\bar{\theta}_i \leq \bar{\theta}_j$, that (in the nontrivial case) the optimal voting-weight vector $\omega \in \Delta N$ has $\omega_i \geq (\bar{\theta}_j/\bar{\theta}_i)\omega_j$.*

EXAMPLE 3 (Uniform value distributions): *Suppose θ_i and θ_j are both uniformly distributed on their respective supports. Direct computation shows φ_i and φ_j are then uniformly distributed on $[2\underline{\theta}_i - \bar{\theta}_i, \bar{\theta}_i]$ and $[2\underline{\theta}_j - \bar{\theta}_j, \bar{\theta}_j]$, respectively. Hence, within this parametric class, the distributional ranking of φ_i and $\alpha\varphi_j$ reduces to $\bar{\theta}_i \leq \alpha\bar{\theta}_j$ and $2\underline{\theta}_i - \bar{\theta}_i \leq \alpha(2\underline{\theta}_j - \bar{\theta}_j)$. Applying Theorem 2 tells us (in the nontrivial case) the optimal voting weight vector $\omega \in \Delta N$ has either $\omega_i = \omega_j = 0$ or $\omega_i > \omega_j$, if $2(\underline{\theta}_j - \underline{\theta}_i) > \bar{\theta}_j - \bar{\theta}_i > 0$.*

dering of their negative value distributions.

¹¹One can easily show the conclusion also holds for *some* optimal voting-weight vector, in the case in which the optimal allocation rule is trivial.

5. Simple Mechanisms

In this section, we study simple mechanisms. We first formulate a permissive class of posted-price mechanisms in which the allocation rule is potentially flexible but the price is fixed, and show a unanimous posted-price mechanism is maximally profitable within this class. Next, we show that in all interesting cases of our model, no optimal mechanism is dominant-strategy incentive compatible. Finally, we compare optimal mechanisms to simple ones as the number of agents grows, showing unanimous posted-pricing profit can perform arbitrarily poorly relative to optimal profit.

5.1. Posted-price mechanisms

An influential result in the mechanism design literature is that a take-it-or-leave-it posted price is the optimal mechanism for selling a single indivisible good to a single agent (Myerson, 1981; Riley and Zeckhauser, 1983). This type of mechanism is ubiquitous and simple and enjoys appealing computational properties. Moreover, beyond the single-agent setting, environments have been identified in which such pricing mechanisms remain approximately optimal (Chawla et al., 2010; Chawla et al., 2015; Hart and Nisan, 2017; Babaioff et al., 2020). A natural question, then, is whether posted-price mechanisms remain optimal for our seller. Having focused on characterizing optimal implementable allocation rules, with relatively little attention paid to the exact implementing transfers, our analysis to this point has left this question unaddressed.

Logically prior to the above question about optimal mechanisms is the question of how one should define a posted-price mechanism. In the one-agent setting, the IC direct mechanisms that correspond to a posted price are exactly 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 the agent freely deciding whether to execute trade—is less immediate to generalize to the multi-agent setting. Who decides whether trade occurs? Once the seller announces a price for the good, a complex negotiation process could ensue between the agents in deciding whether to buy. Might eventual trade outcomes arise from a mixed-strategy equilibrium of the resulting bargaining game between the agents? Can the seller intervene and aid the bargaining process?

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.¹² We also introduce a specific, interpretable pricing mechanism that

¹²As will be clear, any more restrictive definition of a collective posted price would leave the

will be important for our results.

DEFINITION 8: *A mechanism (x, m) is a **collective posted-price** mechanism if some $p \in \mathbb{R}_+$ exists such that $m = px$. Say it is a **unanimous posted-price** mechanism if it is a collective posted-price mechanism in which $x(\theta) = \mathbb{1}_{\min_{i \in N} \theta_i \geq p}$ for every $\theta \in \Theta$.*

One can envision several examples of collective posted-price mechanisms. For example, the seller could set a price p and execute a sale if and only if all agents agree to the purchase—defined above as a unanimous posted price. Alternatively, the principal could post a price and select an agent, or even a subset of agents, perhaps randomly, and sell the good if all the agents in this chosen subset agree to the purchase. Another mechanism would post the price and execute trade if and only if at least one agent, or perhaps some majority of agents, champions the sale.

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

PROPOSITION 2 (Optimal price is unanimous): *A unanimous posted-price mechanism, with price p that solves*

$$\max_{p \geq c} \left\{ (p - c) \prod_{j \in N} [1 - F_j(p)] \right\},$$

generates the highest profit among all IC and IR collective posted-price mechanisms.

Toward establishing this result, observe that IR implies an agent’s interim allocation is zero whenever the agent’s type is below the price. Therefore, trade has zero probability conditional on *any* agent having a realized valuation below the price. From this observation, it follows that any collective posted-price mechanism is outperformed by some unanimous one. Indeed, if the price were weakly below the production cost c , profit would be nonpositive, and a unanimous posted price of c would be at least as good; and if the price were $p > c$, a unanimous posted price of p would generate profitable trade with a higher probability.

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 they never are, restricted to interesting instances of our model. Specifically, whenever factoring in multiple agents’ information at all is optimal (that is, the optimal allocation rule is non-dictatorial), using their reports to fine-tune the price of trade is optimal.

results of this subsection unchanged, as long as it had unanimous posted pricing as a special case.

PROPOSITION 3 (Posted prices are suboptimal): *The following are equivalent.*

1. *Some optimal mechanism (x, m) is a collective posted-price mechanism.*
2. *An i -dictatorship is optimal for some $i \in N$.*

Because an optimal i -dictatorship mechanism is a posted price (and this price is below other agents' lowest types by Corollary 2), we need only show a non-dictatorship optimal mechanism is not a collective posted price.¹³ Because the never-trade and always-trade mechanisms are special instances of dictatorship mechanisms, we need only focus on nontrivial mechanisms. Moreover, in light of Proposition 2, it suffices to show it is not a unanimous posted price. Thus, consider an agent who receives weight $\omega_i \in (0, 1)$ in the optimal allocation rule. His interim allocation probability is non-constant and varies smoothly with his type, and therefore cannot be a step function. But a unanimous posted price would make this interim allocation rule a step function, so that the two cannot coincide.

5.2. Dominant strategies

The notion of incentive compatibility we have used so far in our analysis is *Bayesian* incentive compatibility (BIC, which we denoted as IC earlier in the paper), which requires only that agents' reports be best responses in expectation, given their own realized types. Similarly, as a participation constraint, we required that each agent (knowing his own type) prefers in expectation to interact with the mechanism rather than walk away. Here, we consider more demanding incentive and participation constraints, which we formalize through direct mechanisms below (in light of the revelation principle).

DEFINITION 9: *Say a mechanism (x, m) is **dominant-strategy incentive compatible (DIC)** if*

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

*and say it is **ex-post individually rational (epIR)** if*

$$\theta_i x(\theta) - m(\theta) \geq 0, \quad \forall i \in N, \quad \forall \theta \in \Theta. \quad (\text{epIR})$$

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. Likewise, the mechanism is epIR if an agent (knowing his own type) would rather interact with the mechanism than take an outside option of zero, even if he knew others' reported types.

We showed in Lemma 2 that for a given allocation rule, interim monotonicity is

¹³One can alternatively derive this result as a direct consequence of Propositions 2 and 4. Because the proof of the latter is somewhat involved, we prefer to include a simpler direct proof.

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 a player-weighted virtual value exceeds the cost of production. 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: in all interesting cases of our model, the restriction to DIC mechanisms is with loss of optimality for the seller. If multiple agents must be consulted, optimal mechanisms must leverage agents' uncertainty about others' realized types.

PROPOSITION 4 (Dominance binds): *The following are equivalent:*

1. *An optimal IC, IR mechanism exists that is also DIC and epIR.*
2. *An optimal IC, IR mechanism exists that is also DIC.*
3. *An i -dictatorship mechanism is optimal for some $i \in N$.*

Because a dictatorship conditions only on one agent's private information, it is immediate that implementing transfers can be chosen (e.g. with a posted price) to ensure DIC and epIR are satisfied. The main content of the proposition, then, is that the second condition implies the third. Suppose the optimal allocation rule x characterized in Theorem 1 is DIC when paired with some transfer rule; we aim to show it is a dictatorship.

To prove this feature, we leverage the fact that the essentially unique 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 ($p \in \mathbb{R}_+$) that will be collected if and only if trade occurs and a subsidy ($s \in \mathbb{R}$) that will be granted to the agents 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 an increasing $\{0, 1\}$ -valued transformation of the set-valued function $\theta \mapsto (\mathbb{1}_{\theta_j \geq p})_{j \in N}$. The “coalitional” property amounts to a more explicit description of such increasing functions.

To prove the optimal allocation must be a dictatorship if it is DIC implementable, we apply the structural lemma. Specifically, let the set of possible virtual value profiles be $Z := \prod_{j \in N} [\varphi_j(\underline{\theta}_j), \varphi_j(\bar{\theta}_j)]$, and the set of profiles at which trade occurs be $Z^* := \bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{z \in Z : z_j \geq \varphi_j(p)\}$. Because x is equivalent to x_ω for some weight vector $\omega \in \Delta N$, it follows that Z^* and $Z \setminus Z^*$ are both convex. However, looking at averages of virtual value profiles near the boundary of Z^* reveals Z^* cannot be convex if \mathcal{J} contains two different minimally sufficient coalitions, and $Z \setminus Z^*$ cannot be convex if some minimal coalition in \mathcal{J} contains at least two agents. It follows that the mechanism is equivalent to either the never-trade mechanism ($\mathcal{J} = \emptyset$), the always-trade mechanism ($\mathcal{J} = \{\emptyset\}$), or a nontrivial i -dictatorship ($\mathcal{J} = \{i\}$) for some $i \in N$. The result follows.

5.3. Many agents and suboptimality of posted pricing

Our model, with each agent needing to approve the mechanism (captured by our IR constraint), is perhaps best suited to model bargaining with a small or moderate group of agents. Nevertheless, because our seller's optimum typically combines the private information of multiple agents in a detailed way, it is natural to study the behavior of our model in the limiting case in which the number of agents is large. Doing so is the goal of the present subsection.

Suppose all agents are identical, with $\underline{\theta}_i = \underline{\theta}_1$, $\bar{\theta}_i = \bar{\theta}_1$, and $F_i = F_1$ for every $i \in N$. Holding fixed these parameters and the cost of production, we can think of our model as being parameterized by the number of agents, $N \in \mathbb{N}$.

The result below characterizes the limiting profit that our seller derives from employing an optimal mechanism, in the many-agent limit. Also, because simple mechanisms seem especially appealing when many agents are collectively deciding, we record the asymptotic profit of using a posted price.

PROPOSITION 5 (Asymptotically optimal mechanisms): *Suppose all agents are identical, with types drawn via F_1 . Then, as $N \rightarrow \infty$:*

1. *If $c \neq \underline{\theta}_1$, then the optimal profit converges to $(\underline{\theta}_1 - c)_+$, and the highest profit among collective posted-price mechanisms converges to the same.*
2. *If $c = \underline{\theta}_1$, then the optimal profit converges to $\max_{p \in \mathbb{R}_+} (p - c)[1 - F_1(p)] > 0$, and the highest profit from a collective posted-price converges to zero.*

To prove the proposition, first note the set of optimal voting weights is convex and invariant to coordinate permutations, and hence, the uniform weight vector is optimal. But then, trade occurs in the optimal allocation rule if and only if the (unweighted) *average* virtual value exceeds the cost. Because agents' virtual values are independent (and bounded), the distribution of their average converges to a degenerate distribution on their expectation, $\mathbb{E}[\varphi_1] = \underline{\theta}_1$. From here, the proposition's first point is straightforward. The probability of trade converges to 1 if $\underline{\theta}_1 > c$ and to 0 if $\underline{\theta}_1 < c$. Hence, in these two special cases, the always-trade mechanism or never-trade mechanism is asymptotically optimal. But both of these asymptotically optimal mechanisms are specific instances of unanimous posted-price mechanisms, and so this restricted class of mechanisms has the same asymptotic payoff.

All that remains is the specific case in which $c = \underline{\theta}_1$. This case is in some sense canonical, capturing the case of revenue maximization with no a priori lower bound on the potential gains from trade. First, consider posted-price mechanisms; by Proposition 2, we can focus on unanimous posted prices. Let p be any asymptotically profit-maximizing posted price, that is, any limit point of profit-maximizing unanimous posted prices as $N \rightarrow \infty$. If $p \leq c$, the asymptotic profit per trade is nonpositive; and if $p > c = \underline{\theta}_1$, the asymptotic probability of trade is $\lim_{N \rightarrow \infty} [1 - F_1(p)]^N = 0$. In either case, positive profit is not attainable in the limit with a posted price.

Now, if $c = \underline{\theta}_1$, let us consider the optimal (IC and IR) mechanism. Our limit characterization relies on the fact that a representative agent (agent 1) faces a very predictable average of *other* agents' virtual values. Therefore, the centered random variable

$$\kappa := \frac{N-1}{N} \left(c - \sum_{j \neq 1} \frac{1}{N-1} \varphi_j \right)$$

is independent of agent 1's private information, and it is very concentrated around zero. Because this concentration is strong enough (as quantified by a Hoeffding bound) and trade occurs if and only if $\frac{1}{N}(\varphi_1 - c) \geq \kappa$, it follows that the allocation probability conditional on θ_1 is very close to $\mathbb{1}_{\varphi_1 \geq c}$. Therefore, the profit of the optimal mechanism approaches the profit from a one-agent problem with a virtual value cutoff of c —exactly the optimal profit of the monopolist problem with one buyer whose value has a CDF of F_1 . The proposition follows.

Specializing Proposition 5 to the case of revenue maximization with no a priori

lower bound on agents' valuations ($c = \underline{\theta}_1 = 0$) is instructive. In this canonical case, unanimous posted pricing yields a limit expected revenue of zero as the number of agents grows, whereas the optimal mechanism yields a strictly positive limit profit. Hence, in the many-agent limit, a restriction to simple pricing mechanisms is unboundedly costly to the seller.

6. Alternative Bargaining Arrangements

In this section, we consider variants of the bargaining arrangement in our main model. First, we study the full range of Pareto-optimal mechanisms and show the analysis we used to derive our seller-optimal mechanism can be applied to understand the entire Pareto frontier. Then, we consider how one might relax the veto bargaining constraint on the agents, and discuss the consequences of various modeling choices for optimal mechanisms. Next, we observe a fixed cost-sharing rule among the members of a buyer group can provide a foundation for our assumption that transfers are collective, and explain how our results on heterogeneity can be applied to study asymmetric cost-sharing rules. Finally, we briefly contrast our results with a more traditional model of public good provision in which agent-specific transfers are permitted.

6.1. Pareto-optimal mechanisms

To this point, we have focused on mechanisms that maximize the seller's expected profit. Although this objective is a natural benchmark, it implicitly assumes the seller has extreme bargaining power relative to the buyer group. In this section, we ask what mechanisms might naturally arise with different allocations of bargaining rights. Specifically, we ask which trade outcomes can arise in a Pareto-optimal mechanism. Our characterization of implementable allocation rules, along with the analytical approach we adopt in developing Theorem 1, proves useful in settling this more general question.

Any mechanism (x, m) generates a vector $v_{x,m}$ of $N + 1$ payoffs: $v_{x,m}^i = \mathbb{E}[\theta_i x(\theta) - m(\theta)]$ for each agent $i \in N$, and $v_{x,m}^{N+1} = \mathbb{E}[m(\theta) - cx(\theta)]$ for the seller. Say a mechanism (x, m) is **Pareto optimal** if it is IC and IR, and any alternative IC and IR mechanism (\tilde{x}, \tilde{m}) with $v_{\tilde{x},\tilde{m}} \geq v_{x,m}$ has $v_{\tilde{x},\tilde{m}} = v_{x,m}$.

In the following theorem, we show how the seller-optimal allocation rule generalizes to other Pareto optima. Whereas profit maximization entails allocating the good if and only if a weighted average of virtual values exceeds the cost, general Pareto optimality requires that the good be allocated if and only if a weighted average of *virtual and true* values exceeds the cost.

THEOREM 3 (Pareto-efficient allocations): *An interim-monotone allocation rule x^* is part of some Pareto-optimal mechanism if and only if some $\gamma \in [0, 1]$ and*

$\lambda, \omega \in \Delta N$ exist such that

$$x^*(\theta) = \mathbb{1}_{(1-\gamma)\omega \cdot \varphi + \gamma\lambda \cdot \theta \geq c}$$

almost surely, and $\text{supp}(\omega) \subseteq \text{argmin}_{i \in N} \mathbb{E} [\varphi_i \mathbb{1}_{(1-\gamma)\omega \cdot \varphi + \gamma\lambda \cdot \theta \geq c}]$.

Toward a proof, we first observe, given any $\gamma \in [0, 1]$ and λ , that an essentially unique interim-monotone allocation rule x maximizes the quantity

$$g^{\gamma, \lambda}(x) := \min_{i \in N} \mathbb{E} \{x(\theta) [(1 - \gamma)\varphi_i + \gamma\lambda \cdot \theta - c]\},$$

and that it takes the form described in the above theorem. Indeed, this result can be proven by following identically the proof of Theorem 1 but modifying the objective in the auxiliary zero-sum game. Given this observation, we need only show Pareto-optimal allocation rules are exactly those that maximize $g^{\gamma, \lambda}$ for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$.

First, suppose x^* is part of a Pareto-optimal mechanism, and let $v^* \in \mathbb{R}^{N+1}$ denote the vector of expected payoffs it generates. Pareto optimality tells us v^* is on the boundary of \hat{V} , the set of payoff vectors weakly below one generated by an IC and IR mechanism. Hence, one can find a supporting hyperplane for \hat{V} at v^* —a nonzero $\tilde{\lambda} \in \mathbb{R}^{N+1}$ such that v^* maximizes $\tilde{\lambda} \cdot v$ over all $v \in \hat{V}$. That \hat{V} is downward comprehensive immediately implies all Pareto weights are nonnegative, and that uniformly decreasing the transfer preserves IC and IR implies the Pareto weight on the seller is at least as high as the sum of weights on the agents. Rescaling the weight vector if needed, we may write $\tilde{\lambda} = (\gamma\lambda, 1)$ for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$. For any interim-monotone allocation rule x , then, we can construct an implementing transfer rule (as in Lemma 2) with the property that IR binds for some agent; the result is a payoff vector v satisfying $\tilde{\lambda} \cdot v = g^{\gamma, \lambda}(x)$. Hence, that $\tilde{\lambda}$ supports \hat{V} at v^* implies x^* maximizes $g^{\gamma, \lambda}$.

Conversely, suppose x^* maximizes $g^{\gamma, \lambda}$ for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$, and fix an implementing transfer rule in which some agent's IR constraint binds. Then, because we proved (following the uniqueness part of Theorem 1) this maximizer is essentially unique, it follows that any alternative IC and IR mechanism either generates the same payoff vector or generates a strictly lower $(\gamma\lambda, 1)$ -weighted sum of payoffs. Because this latter weight vector has only nonnegative entries, it follows that the alternative mechanism generates a strictly lower payoff to some player unless it yields an identical payoff vector.

6.2. Beyond veto bargaining

An important feature of our environment is that any agent can unilaterally veto the mechanism. This feature is captured by the requirement that the mechanism be individually rational for all of the agents. Whereas many institutions, especially those comprising a small group of parties, assign veto rights to individual actors,

a more permissive bargaining arrangement may be more appropriate for modeling some contexts. For example, rather than unanimity, a buyer group might require that the terms of trade be approved by at least n agents, where the parameter $n \in \{1, \dots, N\}$ could have $n < N$.¹⁴ This flexibility raises new modeling questions concerning how, exactly, one determines whether a mechanism has sufficient approval.

One possible formulation is that at least n agents must agree, independent of their type realizations, to interact with the seller. This modeling choice might be appropriate, for instance, if we interpret the relationship between the buyer group and seller as an ongoing one, whereas the payoff shocks are idiosyncratic to a particular good or product. 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 seller 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 a weighted sum of the chosen n agents' virtual values exceeds the production cost. The seller would then choose to tailor the mechanism to the n agents she finds most favorable to interact with ex-ante—for instance, the n agents with the highest virtual value distributions if these distributions are first-order stochastically ranked.¹⁵

However, considering the case in which the seller can engage in “coalition building” is also reasonable. That is, the seller could gather support from agents based on their realized willingness to pay for the good at the interim stage. Observe, however, that formalizing a specific protocol to meaningfully capture such a seller's constraints is not straightforward. For example, suppose the seller could require that the agents submit “yes” or “no” votes, and she could commit to the terms of trade as long as at least n agents vote yes. When $n < N$, this ostensibly natural model enables the seller to extract arbitrarily large profits by offering a mechanism that collects transfer $m \gg 0$ and never provides the good. Following standard reasoning, if all agents vote yes, no agent is pivotal, resulting in an equilibrium. Moreover, the mechanism can be modified to make such a voting outcome trembling-hand perfect (hence, not weakly dominated) by setting the transfer equal to $m + 1$ if the vote passes non-unanimously. Although this view is subjective, we feel such constructions sidestep the strategic tradeoffs that should inform the optimal design of selling mechanisms to a buyer group with heterogeneous preferences.

Toward relaxing the veto-bargaining constraint while retaining the strategic

¹⁴A more general bargaining structure, allowing for asymmetry in agents' veto rights, would specify a nonempty $\mathcal{J} \subseteq 2^{N \setminus \{\emptyset\}}$ and require that the mechanism be approved by all $j \in J$ for at least one $J \in \mathcal{J}$.

¹⁵If the virtual values are even ranked according to the hazard-rate order, then, following Theorem 2, the agent with the n^{th} highest distribution would have the highest voting weight.

tradeoffs of our prior analysis, consider the following ad hoc protocol. Agents (having observed their own types) first submit a yes or no vote to a mediator. The mediator informs the seller and agents whether the vote has passed, that is, whether at least n agents voted yes. If the vote passes, the seller’s mechanism—an allocation and a transfer as a function of messages the agents send—is implemented. Notice the seller does not observe the profile of votes in this formulation, and so cannot extract arbitrarily large revenue with the construction described above. However, analyzing this protocol is potentially challenging, requiring substantively different analytical techniques, for another reason. When an agent learns the vote has passed, his conditional belief about other agents’ types depends on how he voted, bringing us outside of the independent private values setting. Agents’ reporting incentives will generally depend on their conditional beliefs about others’ values (which will exhibit correlation), and therefore on their beliefs about other agents’ beliefs, and so on.

The broader point of the above discussion is that we know of no canonical, tractable bargaining framework that relaxes the veto-bargaining constraint while preserving the spirit of our framework. Nevertheless, one may reasonably expect that any such framework would allow the seller to employ the following “posted-price” mechanism: the seller posts a price p and asks agents to vote “yes” or “no.” If and only if at least n agents are in favor, the good is sold at a price p . Let us observe this mechanism can, in some circumstances, perform strictly better than if the seller had no ability to pursue contingent coalition building. For instance, suppose $n = 1$, all agents have the same highest possible value, and the production cost is strictly lower than this highest value so that positive profit is possible. In this case, a best mechanism among those guaranteeing IR for some fixed agent is a posted price of p for the consulted agent. Clearly, the seller attains a strictly higher profit by setting the same price and selling if and only if at least one of the N agents wants to buy.

As the above example illustrates, even simple mechanisms can sometimes yield a higher profit for the seller when she does not face the stringent veto-bargaining constraint. However, formulating a satisfactory framework to capture richer bargaining environments and possible forms of coalition building within a buyer group seems nontrivial. We view the development and analysis of such models as an exciting avenue for future research.

6.3. Fixed-share payment rules

In our model, agents derive heterogeneous private benefits from the buyer group receiving the good, but the cost of transfers is experienced jointly. Our leading interpretation is that the agents are jointly making decisions about how to spend an organization’s funds, and that the variable θ_i represents agent i ’s privately known marginal rate of substitution between the seller’s good and the organization’s money.

An alternative payoff setting is one in which each agent pays private funds to the seller, but mechanisms are restricted—for unmodeled institutional reasons—to follow a fixed cost-sharing rule. Specifically, one could consider a contracting environment in which a physical outcome consists of a probability $\mathbf{x} \in [0, 1]$ with which the good is sold to the buyer group and a vector $\vec{\mathbf{m}} \in \mathbb{R}^N$ of (signed) transfers, where \mathbf{m}_i is paid to the seller by agent i ; and $\sigma \in \Delta N$ is a fixed sharing rule, with all entries strictly positive, such that only transfer vectors $\vec{\mathbf{m}}$ with $\mathbf{m}_i = \sigma_i \sum_{j \in N} \mathbf{m}_j$ are permitted.¹⁶ For example, such a contracting environment might describe a condominium complex interacting with a maintenance company, with a fixed cost-sharing rule prespecified by the homeowner association’s rules.

Inspection of the seller’s and agents’ objective functions shows the above payoff environment is, up to rescaling agent values, equivalent to the model we have studied in our main analysis. If $(\mathbf{v}_i)_{i \in N}$ are the agents’ respective values for the good being provided, and we assume $\{\mathbf{v}_i\}_{i \in N}$ are independent with regular distributions, all of our reported results apply readily when agents’ types are reparameterized as $\theta_i = \mathbf{v}_i / \sigma_i$ —agent i ’s marginal rate of substitution between the seller’s good and the buyer group’s total expenditure.

As a special case, consider fixed cost sharing when agents’ values $\{\mathbf{v}_i\}_{i \in N}$ are identically distributed, and assume without loss that $\sigma_1 \leq \dots \leq \sigma_N$. Letting \tilde{F} and \tilde{f} denote the cumulative distribution function and density, respectively, for \mathbf{v}_i , we can define the scaled virtual value $\psi_i := \mathbf{v}_i - \frac{1 - \tilde{F}(\mathbf{v}_i)}{\tilde{f}(\mathbf{v}_i)}$. Direct computation shows agent i ’s virtual value is $\varphi_i = \bar{\theta}_i \psi_i$, so that \mathbf{v}_i is an increasing transformation of ψ_i . Following Theorem 1, the optimal mechanism allocates the good if and only if $\sum_{i \in N} (\omega_i \bar{\theta}_i) \psi_i$ exceeds the cost for some $\omega \in \Delta N$, and this weight vector is unique in the nontrivial case.

If the cost-sharing rule σ is asymmetric—that is, $\sigma \neq (1/N, \dots, 1/N)$ —the random variables $\{\theta_i\}_{i \in N}$ have proportional but non-identical distributions. Consequently, our results on agent heterogeneity (e.g., Theorem 2) can be directly applied to assess the affects of heterogeneity of the cost-sharing rule. Theorem 2 tells us $\omega_1 \bar{\theta}_1 \geq \dots \geq \omega_N \bar{\theta}_N$, and our leading example illustrates that these inequalities may be strict. Hence, the endogenous voting weights are ranked in accordance with the exogenous cost-sharing weights, with agents that bear a greater fraction of the cost having greater influence over realized trade decisions.

6.4. Agent-specific payments

A substantive restriction of our environment is that our seller cannot collect individual payments from separate agents. In this section, we briefly contrast our setting with that of Güth and Hellwig (1986), in which a profit-maximizing seller offers a mechanism to provide a public good with agent-specific transfers. Of

¹⁶If some share σ_i were zero, the optimal mechanism would ignore agent i ; hence, we omit this possibility for notational simplicity.

course, the models with and without separable transfers are appropriate for studying different applications. To see that the difference affects the formal mechanism design setting as well, let us now observe two different ways in which the results differ across these two settings.

First, optimal mechanisms enjoy different robustness properties. Indeed, the equivalence result of [Gershkov et al. \(2013\)](#) tells us an equivalence between BIC implementability and DIC implementability is a general feature of mechanism design settings with unidimensional independent types, linear utilities, and agent-specific transfers. This result applies directly to the setting of [Güth and Hellwig \(1986\)](#), so that implementation in dominant strategies is without loss of optimality for their seller. By contrast, our results show dominant-strategy implementation is with loss of optimality in every interesting instance of our model ([Proposition 4](#)). Hence, a restriction to collective transfers substantially changes the available robustness properties of optimal mechanisms, even in the canonical setting of the provision of a single good.

Even restricting attention to Bayesian incentive compatibility—for which, given [Lemma 2](#), collective transfers impose no restrictions on implementable allocation rules—allowing for agent-specific transfers substantively changes the form of optimal mechanism.

To contrast the two allocation rules, focusing on a special case of the model, namely that of proportional distributions ([Example 2](#)), is instructive. Suppose $\bar{\theta}_1 < \dots < \bar{\theta}_N$, and suppose the agents' scaled valuations $\mathbf{v}_i := \theta_i/\bar{\theta}_i$ are identically distributed. Thus, the agents are labeled from lowest to highest in terms of their value distributions. Constructing ψ_i from \mathbf{v}_i exactly as in the previous subsection, the variables $\{\psi_i\}_i$ are identically distributed, with θ_i being an increasing transformation of ψ_i for each $i \in N$.

With agent-specific transfers, the good is optimally allocated if and only if $\sum_{i \in N} \varphi_i = \sum_{i \in N} \bar{\theta}_i \psi_i$ exceeds the production cost (see [Güth and Hellwig, 1986](#), [Proposition 4.3](#)). By contrast, in our collective-transfers model, it is allocated if and only if $\sum_{i \in N} (\omega_i \bar{\theta}_i) \psi_i$ exceeds the cost for some $\omega \in \Delta N$. Given that $\bar{\theta}_1 < \dots < \bar{\theta}_N$, [Theorem 2](#) tells us $\omega_1 \bar{\theta}_1 \geq \dots \geq \omega_N \bar{\theta}_N$. Thus, our optimal mechanism overweights agents with lower value distributions, whereas admitting agent-specific transfers induces the opposite ranking for agents' relative influence. Hence, a restriction to collective transfers substantively changes the allocative properties of optimal mechanisms.

References

- S. Alaei, H. Fu, N. Haghpanah, J. Hartline, and A. Malekian. Efficient computation of optimal auctions via reduced forms. *Mathematics of Operations Research*, 44(3):1058–1086, 2019.

- Y. Azrieli and S. Kim. Pareto efficiency and weighted majority rules. *International Economic Review*, 55(4):1067–1088, 2014.
- 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.
- K. C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica: Journal of the Econometric Society*, pages 1175–1187, 1991.
- 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.
- Y. Cai, C. Daskalakis, and S. M. Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 459–478, 2012.
- 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.
- S. Chawla, D. Malec, and B. Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- Y.-K. Che, J. Kim, and K. Mierendorff. Generalized reduced-form auctions: A network-flow approach. *Econometrica*, 81(6):2487–2520, 2013.
- P. Cramton, R. Gibbons, and P. Klemperer. Dissolving a partnership efficiently. *Econometrica: Journal of the Econometric Society*, pages 615–632, 1987.
- C. d’Aspremont and L.-A. Gérard-Varet. Incentives and incomplete information. *Journal of Public economics*, 11(1):25–45, 1979.
- C. d’Aspremont, J. Crémer, and L.-A. Gérard-Varet. Balanced bayesian mechanisms. *Journal of Economic Theory*, 115(2):385–396, 2004.
- 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.
- A. Gershkov, B. Moldovanu, and X. Shi. Optimal voting rules. *The Review of Economic Studies*, 84(2):688–717, 2017.
- 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.
- W. Güth and M. Hellwig. The private supply of a public good. *Journal of Economics*, 46(1):121–159, 1986.

- S. Hart and N. Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.
- P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of economic theory*, 85(2):258–293, 1999.
- R. Kirkegaard. A mechanism design approach to ranking asymmetric auctions. *Econometrica*, 80(5):2349–2364, 2012.
- V. Krishna and J. Morgan. Majority rule and utilitarian welfare. *American Economic Journal: Microeconomics*, 7(4):339–75, 2015.
- G. J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents. *The Review of Economic Studies*, 57(3):351–367, 1990.
- L. Makowski and C. Mezzetti. Bayesian and weakly robust first best mechanisms: characterizations. *Journal of Economic Theory*, 64(2):500–519, 1994.
- A. M. Manelli and D. R. Vincent. Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 78(6):1905–1938, 2010.
- A. M. Manelli and D. R. Vincent. Dominant-strategy and bayesian incentive compatibility in multi-object trading environments. *Journal of Mathematical Economics*, 82:214–226, 2019.
- E. Maskin and J. Riley. Asymmetric auctions. *The review of economic studies*, 67(3):413–438, 2000.
- R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- M. J. Osborne and A. Rubinstein. *A course in game theory*. MIT press, 1994.
- D. W. Rae. Decision-rules and individual values in constitutional choice. *The American Political Science Review*, 63(1):40–56, 1969.
- 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.
- R. Rob. Pollution claim settlements under private information. *Journal of Economic Theory*, 47(2):307–333, 1989.
- P. W. Schmitz and T. Tröger. The (sub-) optimality of the majority rule. *Games and Economic Behavior*, 74(2):651–665, 2012.
- M. Shaked and J. G. Shanthikumar. *Stochastic orders*. Springer Science & Business Media, 2007.
- A. C.-C. Yao. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 3–20, 2017.

A. Appendix: Proofs

A.1. Proofs for Section 3

Proof of Lemma 1. If transfer rule m has $M_i^m = M_i^*$ for each $i \in N$, then iterated expectations implies

$$\mathbb{E}[M_i^*(\boldsymbol{\theta}_i)] = \mathbb{E}[m(\boldsymbol{\theta})] = \mathbb{E}[M_j^*(\boldsymbol{\theta}_j)] \quad \forall i, j \in N.$$

Conversely, suppose $\bar{m} \in \mathbb{R}$ is such that $\mathbb{E}[M_i^*(\boldsymbol{\theta}_i)] = \bar{m}$ for every $i \in N$. It then follows immediately from independence of $\{\boldsymbol{\theta}_i\}_{i \in N}$ that the following transfer rule generates the desired interim versions for any constant $\mathbf{m}_0 \in \mathbb{R} \setminus \{\bar{m}\}$:

$$\begin{aligned} m : \Theta &\rightarrow \mathbb{R} \\ \theta &\mapsto \mathbf{m}_0 + \frac{1}{(\bar{m} - \mathbf{m}_0)^{N-1}} \prod_{i \in N} [M_i^*(\theta_i) - \mathbf{m}_0]. \end{aligned}$$

Now, we turn to the final assertion. If nonnegative transfer rule m has $M_i^m = M_i^*$ for $i \in N$, then monotonicity of integration implies $M_i^*(\theta_i) \geq 0$ for every $\theta_i \in \Theta_i$. Conversely, suppose each of $\{M_i^*\}_{i \in N}$ is nonnegative, and their expectations are all equal to $\bar{m} \in \mathbb{R}$. Monotonicity of integration then implies $\bar{m} \geq 0$. If $\bar{m} > 0$ [resp. $\bar{m} = 0$], then the above-constructed transfer rule is nonnegative when for $\mathbf{m}_0 = 0$ [resp. $\mathbf{m}_0 = -1$].¹⁷ \square

Proof of Lemma 2. Let $X_i := X_i^x$ for each $i \in N$. Given a transfer rule m , standard arguments (Myerson, 1981) show that (x, m) is IC if and only if each $i \in N$ has X_i weakly increasing and

$$M_i^m(\theta_i) = X_i(\theta_i)\theta_i - \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i) d\tilde{\theta}_i - \underline{U}_i, \quad \forall \theta_i \in \Theta_i \quad (1)$$

for some constant $\underline{U}_i \in \mathbb{R}$; that such a mechanism is IR if and only if $\underline{U}_i \geq 0$ for each $i \in N$; and that any $M_i : \Theta_i \rightarrow \mathbb{R}$ satisfying equation (1) has $\mathbb{E}[M_i(\boldsymbol{\theta}_i)] = \mathbb{E}[X_i(\boldsymbol{\theta}_i)\boldsymbol{\varphi}_i] - \underline{U}_i$. The latter expression implies that, by iterated expectations, any transfer rule m such that (x, m) is IC has

$$\mathbb{E}[m(\boldsymbol{\theta})] = \mathbb{E}[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i] - \underline{U}_i \quad \forall i \in N. \quad (2)$$

Let us now observe how the two parts of the lemma follow from the above standard observations. For the necessity of interim-monotonicity in the first part, nothing remains to show. To see that the payoff expression in the second part is an upper bound on attainable profits, note that every IR and IC mechanism (x, m) generates, by equation (2), a profit of

$$\Pi(x, m) = \min_{i \in N} \{\mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)] - \underline{U}_i\} \leq \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)].$$

¹⁷We should note that our proof does not use any of our regularity assumptions on the type spaces and their distributions: Types can be independently drawn from any probability spaces whatsoever. Moreover, boundedness can be replaced with integrability, in which case the implementing m need not be bounded. Further, under an appropriate absolute continuity condition, independence is also dispensable for this result.

So the lemma will follow if we can construct, given an arbitrary allocation rule x whose induced interim allocation rules $\{X_i\}_{i \in N}$ are all weakly increasing, a transfer rule m such that (x, m) is IC and IR with $\Pi(x, m) = \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)]$. To that end, fix some $i_* \in \operatorname{argmin}_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})\boldsymbol{\varphi}_i]$. For each $i \in N$, define the hypothetical payoff lower bound $\underline{U}_i := \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - \boldsymbol{\varphi}_{i_*})]$ and hypothetical interim transfer rule

$$\begin{aligned} M_i^* : \Theta_i &\rightarrow \mathbb{R} \\ \theta_i &\mapsto X_i(\theta_i)\theta_i - \int_{\underline{\theta}_i}^{\theta_i} X_i(\tilde{\theta}_i) d\tilde{\theta}_i - \underline{U}_i. \end{aligned}$$

Observe every $i \in N$ has

$$\mathbb{E}[M_i^*(\boldsymbol{\theta}_i)] = \mathbb{E}[x_i(\boldsymbol{\theta}_i)\boldsymbol{\varphi}_i] - \underline{U}_i = \mathbb{E}[x_{i_*}(\boldsymbol{\theta}_{i_*})\boldsymbol{\varphi}_{i_*}] = \mathbb{E}[M_{i_*}^*(\boldsymbol{\theta}_{i_*})],$$

and so Lemma 1 delivers a transfer rule m whose interim transfer rules satisfy $M_i^m = M_i^* \forall i \in N$. Let us now observe (x, m) is as desired. Indeed, interim-monotonicity of x and equation (1) implies IC; that $\{\underline{U}_i\}_{i \in N}$ are all non-negative implies IR; and that $\underline{U}_{i_*} = 0$ implies $\Pi(x, m) = \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_{i_*} - c)] = \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)]$. The lemma follows. \square

Proof of Theorem 1. Let $\tilde{\mathcal{X}}$ denote the set \mathcal{X} , modulo the F -almost everywhere equivalence relation. One can view $\tilde{\mathcal{X}}$ as subset of $L^\infty(\Theta, F)$, and the Banach-Alaoglu theorem then implies $\tilde{\mathcal{X}}$ is weak*-compact.¹⁸ Consider the optimization problem,

$$\max_{x \in \tilde{\mathcal{X}}} \left\{ \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)] \right\}, \quad (\text{RSP})$$

which is our seller's problem without the monotonicity constraint. In what follows, we will pursue a solution to this relaxed problem. As we will show, this program is solved by a unique $x^* \in \tilde{\mathcal{X}}$, and this x^* happens to exhibit monotone interim allocation probabilities. Hence, it will follow that x^* is the unique solution to our seller's problem.

Toward solving (RSP), 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 $\mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)]$. Because $\tilde{\mathcal{X}}$ is weak*-compact and convex (the space ΔN obviously is as well), and the objective as weak*-continuous in the strategy profile, it follows from Sion's minimax theorem that

$$\max_{x \in \tilde{\mathcal{X}}} \min_{\omega \in \Delta N} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)] = \min_{\omega \in \Delta N} \max_{x \in \tilde{\mathcal{X}}} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)],$$

where all maxima/minima in the equation are attained by Berge's theorem.

Because the auxiliary game is zero-sum (Proposition 22.2, Osborne and Rubinstein, 1994), the Nash equilibria are exactly the pairs $(x^*, \omega^*) \in \tilde{\mathcal{X}} \times \Delta N$ for

¹⁸As is standard, $L^\infty(\Theta, F)$ is isometrically isomorphic to the dual of $L^1(\Theta, F)$.

which

$$\begin{aligned} x^* &\in \operatorname{argmax}_{x \in \tilde{\mathcal{X}}} \min_{\omega \in \Delta N} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)] \text{ and} \\ \omega^* &\in \operatorname{argmin}_{\omega \in \Delta N} \max_{x \in \tilde{\mathcal{X}}} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)]. \end{aligned}$$

Observe, though, that $\min_{\omega \in \Delta N} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)] = \min_{i \in N} \mathbb{E}[x(\boldsymbol{\theta})(\varphi_i - c)]$ for each $x \in \tilde{\mathcal{X}}$. Hence, x^* maximizes this quantity if and only if x^* solves (RSP). Moreover, $\max_{x \in \tilde{\mathcal{X}}} \mathbb{E}[x(\boldsymbol{\theta})(\omega \cdot \boldsymbol{\varphi} - c)] = \max_{x \in \tilde{\mathcal{X}}} \mathbb{E}[(\omega \cdot \boldsymbol{\varphi} - c)_+]$ for each $\omega \in \Delta N$, so that maximizing the two expressions is equivalent. Finally, because these maxima/minima are obtained, some Nash equilibrium exists.

To summarize our progress so far, we know that a Nash equilibrium exists for the zero-sum game, and Nash equilibria are exactly the pairs $(x^*, \omega^*) \in \tilde{\mathcal{X}} \times \Delta N$ for which x^* solves (RSP) and ω^* solves $\min_{\omega \in \Delta N} \max_{x \in \tilde{\mathcal{X}}} \mathbb{E}[(\omega \cdot \boldsymbol{\varphi} - c)_+]$.

Now, for an arbitrary $\omega \in \Delta N$. Because $\{\boldsymbol{\theta}_i\}_{i \in N}$ are atomless and independent and $\{\varphi_i\}_{i \in N}$ are all strictly increasing, it follows that $\mathbb{P}\{\omega \cdot \boldsymbol{\varphi} = c\} = 0$, so that the ω -voting rule x_ω is Minimizer's unique best response to ω . From the product structure of the set of Nash equilibria, then, it follows that Maximizer has a unique Nash equilibrium strategy x^* , which is then the unique solution to (RSP). Moreover, because a voting rule is obviously interim-monotone (given independent types and increasing virtual values), it follows that the unique solution to (RSP) is also the unique solution to the seller's problem (SP).

All that remains for our characterization of optimal allocation rules is to show the equivalence of the two conditions in the theorem's statement for a given $\omega \in \Delta N$, and that these conditions imply x_ω is optimal. We have argued above that the first condition is equivalent to ω being a Nash equilibrium strategy for Minimizer. Meanwhile, because we have argued $x_\omega = \mathbb{1}_{\omega \cdot \boldsymbol{\varphi} \geq c}$ is a unique Maximizer best response to ω , it follows readily that the second condition is equivalent to ω being a Nash equilibrium strategy for Minimizer. Hence the first and second conditions are equivalent. Moreover, we have argued that, if ω is a Nash equilibrium strategy for Minimizer, then the ω -voting rule is an optimal allocation rule. Therefore, if ω satisfies (1) or (2), then x_ω is an optimal allocation rule.

We now address the theorem's final statement—that transfers are without loss taken to be nonnegative. If the never-trade mechanism is optimal, there is nothing to show. So focus on the complementary case in which an optimal mechanism generates strictly positive revenue.

We have proved that the optimal allocation rule is $x = x_\omega$ for some optimal $\omega \in \Delta N$. Let $\{M_i^*\}_{i \in N}$ denote interim transfer rules, constructed in the proof of Lemma 2, that implement x at maximum possible profit. By Lemma 1, it suffices to show M_i^* is nonnegative for each $i \in N$. To see this feature, note that (given the functional form of its construction) M_i^* is always weakly increasing and is constant if X_i^x is constant. With this observation, we can establish that $M_i^* \geq 0$ in two exhaustive cases. First, if $\omega_i = 0$, then X_i^x is constant and so M_i^* is constant, hence equal to $\mathbb{E}[m(\boldsymbol{\theta})] \geq c\mathbb{E}[x(\boldsymbol{\theta})] \geq 0$, where the first inequality holds because an optimal mechanism is weakly better for the seller than the never-trade

mechanism. Second, if $\omega_i > 0$, optimality of ω implies $i \in \operatorname{argmin}_{j \in N} \mathbb{E} [\varphi_j x(\boldsymbol{\theta})]$. But, in this case the constructed transfer rules satisfy

$$0 = \underline{\theta}_i X_i^x(\underline{\theta}_i) - M_i^*(\underline{\theta}_i) \geq -M_i^*(\underline{\theta}_i),$$

so that $M_i^* \geq 0$. \square

Proof of Proposition 1. First, it is immediate that the never-trade [resp. always-trade] mechanism is optimal among all IC and IR mechanisms using allocation rule $x = 0$ [resp. $x = 1$]. Hence, for the first two items, it suffices to show the given condition characterizes when $x = 0$ [resp. $x = 1$] solves (SP).

Now, we characterize when the never-trade allocation rule is optimal. If $\bar{\theta}_i \leq c$ for some $i \in N$, then choosing $\omega \in \Delta N$ with $\omega_j = \mathbb{1}_{i=j}$ for each $j \in N$ satisfies the first condition in Theorem 1, so that the zero allocation rule x_ω is optimal. Conversely, suppose $\bar{\theta}_i > c$ for every $i \in N$. Theorem 1 says some optimal $\omega \in \Delta N$ exists, and the ω -voting rule is a uniquely optimal allocation rule—but observe this rule entails a positive trade probability for any $\omega \in \Delta N$.

Next, we characterize when the always-trade allocation rule is optimal. By Theorem 1, this allocation rule is optimal if and only if some $\omega \in \Delta N$ exists such that x_ω is the always-trade allocation rule and $\operatorname{supp}(\omega) \subseteq \operatorname{argmin}_{i \in N} \mathbb{E} [\varphi_i \mathbb{1}_{\omega \cdot \varphi \geq c}]$. If x_ω is the always-trade allocation rule, though, observe that

$$\mathbb{E} [\varphi_i \mathbb{1}_{\omega \cdot \varphi \geq c}] = \mathbb{E} [\varphi_i] = \underline{\theta}_i.$$

Moreover, for a given $\omega \in \Delta N$, note that x_ω is the always-trade allocation rule if and only if $\omega \cdot \varphi(\underline{\theta}) \geq c$. Hence, the always-trade allocation rule is optimal if and only if some $\omega \in \Delta(\operatorname{argmin}_{j \in N} \underline{\theta}_j)$ exists such that $\omega \cdot \varphi(\underline{\theta}) \geq c$. Because some degenerate such ω maximizes $\omega \cdot \varphi(\underline{\theta})$, this property is equivalent to some $i \in \operatorname{argmin}_{j \in N} \underline{\theta}_j$ existing such that $\varphi_i(\underline{\theta}_i) \geq c$.

Finally, we turn to uniqueness. Suppose $\omega, \tilde{\omega} \in \Delta N$ are such that x_ω and $x_{\tilde{\omega}}$ are both optimal. Theorem 1 shows that $x_{\tilde{\omega}}(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely. Our aim is to show, assuming these allocation rules are nontrivial, that $\omega = \tilde{\omega}$. Toward establishing this equality, define $G := \prod_{i \in N} (\varphi_i(\underline{\theta}_i) - c, \bar{\theta}_i - c)$, the interior of the support of $\varphi - c \mathbb{1}_N$. Now, define the linear map $L : \mathbb{R}^N \rightarrow \mathbb{R}^2$ 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 x_ω and $x_{\tilde{\omega}}$ are nontrivial 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 $\mathbb{P}\{x_{\tilde{\omega}}(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})\} = 1$ 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 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 $\omega = \tilde{\omega}$. \square

A.2. Proofs for Section 4

Proof of Corollary 1. The stochastic dominance hypothesis implies that $\bar{\theta}_i \leq \bar{\theta}_j$ and $\underline{\theta}_i \leq \underline{\theta}_j$ for each $j \in N$. Moreover, it implies that any $j \in N$ with $\underline{\theta}_j = \underline{\theta}_i$

has $f_i(\underline{\theta}_i) \geq f_j(\underline{\theta}_i)$, so that $\varphi_i(\underline{\theta}_i) \geq \varphi_i(\underline{\theta}_j)$. The result then follows directly from Proposition 1. \square

Proof of Corollary 2. The second and third conditions are equivalent, because the profit maximizing price for agent i induces allocation rule $\theta \mapsto \mathbb{1}_{\varphi_i(\theta_i) \geq c}$, and $\mathbb{E}[\varphi_j] = \underline{\theta}_j$ for each $j \in N$. We now turn to showing the first and third conditions are equivalent.

By Theorem 1, this ω is optimal if and only if i minimizes $\mathbb{E}[\varphi_j \mathbb{1}_{\omega \cdot \varphi \geq c}]$ over all $j \in N$. But observe each $j \in N$ has $\mathbb{E}[\varphi_j \mathbb{1}_{\omega \cdot \varphi \geq c}] = \mathbb{E}[\varphi_j \mathbb{1}_{\varphi_i \geq c}]$, which is (because $\{\varphi_j\}_{j \in N}$ are independent) equal to $\mathbb{E}[\varphi_j] \mathbb{P}\{\varphi_i \geq c\}$ if $j \neq i$. Hence ω is optimal if and only if each $j \in N \setminus \{i\}$ has $\mathbb{E}[\varphi_j] \mathbb{P}\{\varphi_i \geq c\} \geq \mathbb{E}[\varphi_i \mathbb{1}_{\varphi_i \geq c}]$.

Consider now two exhaustive cases. First, if $\bar{\theta}_i \leq c$, then $\mathbb{P}\{\varphi_i \geq c\} = \mathbb{E}[\varphi_i \mathbb{1}_{\varphi_i \geq c}] = 0$, and so the inequalities are trivially satisfied. Second, if $\bar{\theta}_i > c$, then dividing the inequalities by $\mathbb{P}\{\varphi_i \geq c\} > 0$ tells us ω is optimal if and only if each $i \in N \setminus \{i\}$ has $\mathbb{E}[\varphi_j] \geq \mathbb{E}[\varphi_i | \varphi_i \geq c]$. \square

Proof of Corollary 3. Consider any $\omega \in \Delta N$ such that x_ω is nontrivial and has $\omega_i = 0$. It suffices to show that ω cannot be optimal.

Fixing some $j \in \text{supp}(\omega)$, observe that

$$\begin{aligned} \mathbb{E}[\varphi_i | \omega \cdot \varphi \geq c] &= \mathbb{E}[\varphi_i] \\ &= \underline{\theta}_i \\ &\leq \underline{\theta}_j \\ &= \mathbb{E}[\varphi_j] \\ &< \mathbb{E}\left[\varphi_j \mid \varphi_j \geq \frac{1}{\omega_j} \left(c - \sum_{k \in N \setminus \{j\}} \omega_k \varphi_k\right)\right] \\ &= \mathbb{E}[\varphi_j | \omega \cdot \varphi \geq c], \end{aligned}$$

where the first equality follows from $\{\varphi_k\}_{k \in N}$ being independent, and the strict inequality follows from the same and from the conditioning event having interior probability. Hence, ω is not optimal. \square

Proof of Theorem 2. Assume for a contradiction that the unique (by Proposition 1) optimal vector of voting weights ω satisfies $\omega_i < \frac{1}{\alpha} \omega_j$. Let

$$\beta := \frac{\alpha(\omega_i + \omega_j)}{\alpha^2 \omega_i + \omega_j} \in (0, 1],$$

and define $\tilde{\omega} \in \mathbb{R}^N$ by letting $\tilde{\omega}_i := \frac{\beta}{\alpha} \omega_j$ and $\tilde{\omega}_j := \beta \alpha \omega_i$, and letting $\tilde{\omega}_k = \omega_k$ for every other $k \in N$. Notice that $\tilde{\omega} \in \Delta N$ since $\tilde{\omega}_i, \tilde{\omega}_j \geq 0$ and

$$\tilde{\omega}_i + \tilde{\omega}_j = \beta \left(\frac{\omega_j}{\alpha} + \alpha \omega_i \right) = \frac{\alpha(\omega_i + \omega_j)}{\alpha^2 \omega_i + \omega_j} \left(\frac{\alpha^2 \omega_i + \omega_j}{\alpha} \right) = \omega_i + \omega_j.$$

Also, note that $\tilde{\omega} \neq \omega$. Indeed, if $\omega_i = 0$ then this fact follows from $\tilde{\omega}_i = \frac{\beta}{\alpha} \omega_j > 0$; and otherwise, it follows from $\frac{\tilde{\omega}_j}{\tilde{\omega}_i} = \alpha^2 \frac{\omega_i}{\omega_j} < \frac{\omega_j}{\omega_i}$. Hence, it suffices to show $\mathbb{E}[(\omega \cdot \varphi - c)_+] \geq \mathbb{E}[(\tilde{\omega} \cdot \varphi - c)_+]$. Indeed, if we could show this ranking, then $\tilde{\omega}$ would be optimal too—in contradiction to the unique optimality of ω .

So let us turn to showing that $\mathbb{E}[(\omega \cdot \boldsymbol{\varphi} - c)_+] \geq \mathbb{E}[(\tilde{\omega} \cdot \boldsymbol{\varphi} - c)_+]$. To prove the result, we invoke results from [Shaked and Shanthikumar \(2007\)](#). First, as $\boldsymbol{\varphi}_i$ is smaller than $\alpha\boldsymbol{\varphi}_j$ in the hazard rate order, it follows from Theorem 1.B.41 that $-\alpha\boldsymbol{\varphi}_j$ is smaller than $-\boldsymbol{\varphi}_i$ in the reverse hazard order. Hence, Theorem 4.A.37 implies that $\omega_i(-\boldsymbol{\varphi}_i) + \frac{1}{\alpha}\omega_j(-\alpha\boldsymbol{\varphi}_j) = -(\omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j)$ is smaller than $\frac{1}{\alpha}\omega_j(-\boldsymbol{\varphi}_i) + \omega_i(-\alpha\boldsymbol{\varphi}_j) = -(\frac{1}{\alpha}\omega_j\boldsymbol{\varphi}_i + \alpha\omega_i\boldsymbol{\varphi}_j)$ in the increasing and concave order. Therefore, by 4.A.1, $\omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j$ is larger than $\frac{1}{\alpha}\omega_j\boldsymbol{\varphi}_i + \alpha\omega_i\boldsymbol{\varphi}_j = \frac{1}{\beta}(\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j)$ in the increasing and convex order. Moreover, $\frac{1}{\beta}(\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j)$ is larger than $\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j$ in the increasing and convex order since $\beta \leq 1$.¹⁹ Therefore, $\mathbb{E}[\eta(\omega_i\boldsymbol{\varphi}_i + \omega_j\boldsymbol{\varphi}_j)] \geq \mathbb{E}[\eta(\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j)]$ for every (weakly) increasing and convex $\eta : \mathbb{R} \rightarrow \mathbb{R}$. The desired inequality then follows from applying this ranking to η given by

$$\eta(y) := \mathbb{E} \left[\left(y - c + \sum_{k \in N \setminus \{i,j\}} \omega_k \boldsymbol{\varphi}_k \right)_+ \right],$$

which is convex and increasing because $(\cdot)_+$ is. □

A.3. Proofs for Section 5

A.3.1. Proofs for Section 5.1

Proof of Proposition 2. Consider an arbitrary collective posted price mechanism (x, m) with price $p \geq 0$. We will show a unanimous posted price performs weakly better.²⁰

If $p \leq c$, then the profit associated with the mechanism is always non-positive, and so a unanimous posted price with price c is weakly better.

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

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

associated with a unanimous posted price of p . Hence, $(p - c)\mathbb{E}[x(\boldsymbol{\theta})] \leq (p - c)\mathbb{E}[x^U(\boldsymbol{\theta})]$ —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. □

Proof of Proposition 3. The equivalence is trivial in the case that the always-trade or never-trade mechanism is optimal. Hence, we restrict attention to the case that the optimal allocation rule is nontrivial; let ω denote the unique (by Proposition 1) optimal voting weights, and $x := x_\omega$. First, if ω is an i -dictatorship for some $i \in N$, then a posted price mechanism with price $p \in \Theta_i$ such that $\varphi_i(p) = c$ is optimal.

¹⁹As the random variable $\mathbf{v} := \frac{1}{\beta}(\tilde{\omega}_i\boldsymbol{\varphi}_i + \tilde{\omega}_j\boldsymbol{\varphi}_j)$ has nonnegative mean, any convex increasing η has $\mathbb{E}\eta(\mathbf{v}) = (1 - \beta)\mathbb{E}\eta(\mathbf{v}) + \beta\mathbb{E}\eta(\mathbf{v}) \geq (1 - \beta)\eta(\mathbb{E}\mathbf{v}) + [\mathbb{E}\eta(\beta\mathbf{v}) - (1 - \beta)\eta(0)] \geq \mathbb{E}\eta(\beta\mathbf{v})$.

²⁰It follows from our proof that, if some IC and IR mechanism generates strictly positive profit, then this payoff ranking can be made strict unless the original mechanism is essentially a unanimous posted price.

Assume now for a contradiction that ω is a dictatorship and some optimal mechanism is a collective posted price mechanism. By Proposition 1, some optimal mechanism is in fact a unanimous posted price mechanism, with price p . Now, let $i \in N$ be such that $0 < \omega_i < 1$ —some such i exists because ω is not a dictatorship.

For each $\theta_i \in \Theta_i$, i 's interim probability of trade is given by

$$X(\theta_i) = \mathbb{P} \left\{ \sum_{j \in N \setminus \{i\}} \omega_j \varphi_j(\boldsymbol{\theta}_j) \geq c - \omega_i \varphi_i(\theta_i) \right\}.$$

Recall that $\boldsymbol{\theta}$ admits a density, ω_{-i} and ω_i are both nonzero, and $\{\varphi_j\}_{j \in N}$ are all continuous and strictly increasing. It follows that the random variable on the left side of the above inequality is atomlessly distributed, while the number on the right side varies continuously with θ_i . Hence, X_i is continuous. Moreover, X_i is not constant, because the allocation rule is nontrivial while the random variable on the left side of the inequality has convex support.

Next, observe that uniqueness of the optimal allocation rule (by Theorem 1) implies $x(\boldsymbol{\theta}) = \mathbb{1}_{\boldsymbol{\theta}_j \geq p \forall j \in N}$ almost surely. Hence, by iterated expectations,

$$X_i(\theta_i) = \bar{x} \mathbb{1}_{\theta_i \geq p} \text{ almost surely,}$$

where $\bar{x} := \mathbb{P} \{ \boldsymbol{\theta}_j \geq p \forall j \in N \setminus \{i\} \} \in [0, 1]$. So the function $X_i : \Theta_i \rightarrow [0, 1]$ is continuous and not constant, and agrees almost everywhere with a $\{0, \bar{x}\}$ -valued function—a contradiction. The proposition follows. \square

A.3.2. Proofs for Section 5.2

LEMMA 3: *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 $x(\theta) = x(\theta') = 1$, and $\theta^* := \theta \wedge \theta'$ if $x(\theta) = x(\theta') = 0$. 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

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

Observe, either $\theta^0 \leq \dots \leq \theta^N$ and $x(\theta^0) = 1$, or $\theta^0 \geq \dots \geq \theta^N$ and $x(\theta^0) = 0$. In either case, because x is weakly increasing (due to DIC) and can only take values in $[0, 1]$, it follows by induction that $x(\theta^0) = \dots = x(\theta^N)$. 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, $m(\theta) = m(\theta^0) = \dots = m(\theta^N) = m(\theta^*)$, as desired. \square

DEFINITION 10: *Say a mechanism (x, m) or an allocation rule x is **bang-bang** if $x(\boldsymbol{\theta}) \in \{0, 1\}$ almost surely.*

LEMMA 4 (DIC mechanisms): *Suppose (x, m) is a DIC bang-bang mechanism. Then, some $p \in \mathbb{R}_+$, $s \in \mathbb{R}$ and $\mathcal{J} \subseteq 2^N$ exist such that, almost surely:*

- (1) $m(\boldsymbol{\theta}) = px(\boldsymbol{\theta}) - s$;
- (2) $x(\boldsymbol{\theta}) = \mathbb{1}_{\cup_{J \in \mathcal{J}} \cap_{j \in J} \{\theta_j \geq p\}}$

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 3, 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]. Moreover, DIC implies $m^L \leq m^H$ if there exist type profiles leading to both allocation probabilities; and we may without loss take $m^L \leq m^H$ in the complementary case. So, defining $p := m^H - m^L \geq 0$ 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 a measure-zero subset of its domain, and similarly modifying the transfer rule to maintain $m = px - s$, we may assume without loss that x is (statewise) $\{0, 1\}$ -valued. Indeed, if $x(\boldsymbol{\theta}) = 0$ almost surely, we can replace the allocation rule with the zero allocation rule; and in the complementary case, we can replace the allocation rule with $\theta \mapsto \mathbb{1}_{x(\theta) > 0}$. It is easy to see that DIC of the modified mechanism follows from DIC of the original one.

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 $(\theta_i - p)x(\hat{\theta}) - s$, 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: $x(\cdot, \theta_{-i}) = 1$, $x(\cdot, \theta_{-i}) = 0$, 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 $\tilde{\Theta} := \prod_{i \in N} [\Theta_i \setminus \{p\}]$, some $y : \{0, 1\}^N \rightarrow \{0, 1\}$ exists such that every $\theta \in \tilde{\Theta}$ has $x(\theta) = y((\mathbb{1}_{\theta_i \geq p})_{i \in N})$. Moreover, we may assume without loss that y is constant in its i coordinate if $p \leq \underline{\theta}_i$ or $p \geq \bar{\theta}_i$ 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, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \tilde{\mathcal{J}}} \bigcap_{j \in J} \{\theta_j \geq p\}}$ 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, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \hat{\mathcal{J}}} \bigcap_{j \in J} \{\theta_j \geq p\}}$ 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, $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{\theta_j \geq p\}}$ 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. \square

Proof of Proposition 4. That the third condition implies the first is immediate: A dictatorship mechanism is trivially DIC because no agent both affects the outcome and has some co-player who affects the outcome. As the first condition obviously implies the second, we need only show the second condition implies the first. To that end, suppose allocation rule x is both optimal and DIC-implementable. Our aim is to show x is a dictatorship.

Theorem 1 implies $x(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely, for some $\omega \in \Delta N$. Lemma 4 implies $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{\theta_j > p\}}$ almost surely, for some $p \in \mathbb{R}$ and $\mathcal{J} \subseteq 2^N$ such that \mathcal{J} are pairwise non-nested and every $j \in N^* := \bigcup \mathcal{J}$ has $\underline{\theta}_j < p < \bar{\theta}_j$.

If we can establish that the set \mathcal{J} is equal to \emptyset , to $\{\emptyset\}$, or to $\{i\}$ for some $i \in N$ —corresponding to never-trade, always-trade, and nontrivial i -dictatorship—then the proposition will follow. Toward establishing this fact, define $Z :=$

$\prod_{i \in N} [\varphi_i(\underline{\theta}_i), \bar{\theta}_i]$, and let $Z^* \subseteq Z$ denote the support of the measure on Z which assigns mass $\mathbb{E} \left[x(\boldsymbol{\theta}) \mathbb{1}_{\varphi \in \hat{Z}} \right]$ to every Borel $\hat{Z} \subseteq Z$. That $x(\boldsymbol{\theta}) = x_\omega(\boldsymbol{\theta})$ almost surely implies both Z^* and $Z \setminus Z^*$ are convex. Meanwhile, that $x(\boldsymbol{\theta}) = \mathbb{1}_{\bigcup_{J \in \mathcal{J}} \bigcap_{j \in J} \{\theta_j > p\}}$ almost surely implies

$$Z^* = \bigcup_{J \in \mathcal{J}} \{z \in Z : z_j \geq \varphi_j(p) \ \forall j \in J\}.$$

Using this characterization, we can show that \mathcal{J} is one of the aforementioned sets.

First, let us see that $|\mathcal{J}| \leq 1$. Assume for contradiction that $J, J' \in \mathcal{J}$ have $J \neq J'$. Define now the elements $z, z' \in Z$ via $z := (\varphi_i(\underline{\theta}_i) \mathbb{1}_{i \notin J} + \varphi_i(p) \mathbb{1}_{i \in J})_{i \in N}$ and $z' := (\varphi_i(\underline{\theta}_i) \mathbb{1}_{i \notin J'} + \varphi_i(p) \mathbb{1}_{i \in J'})_{i \in N}$. Observe that z and z' are in Z^* , but their midpoint is not—contradicting the convexity of Z^* .

Hence, if \mathcal{J} is nonempty, then $\mathcal{J} = \{J\}$ for some $J \subseteq N$. Now, let us see that $|J| \leq 1$ in this case. Assume for a contradiction that $j, j' \in J$ have $j \neq j'$. Given $\epsilon \in (0, \min_{i \in J} [\varphi_i(\bar{\theta}_i) - \varphi_i(p)])$, define $z_\epsilon, z'_\epsilon \in Z$ by letting $z_\epsilon := ([\varphi_i(p) - \epsilon] \mathbb{1}_{i=j'} + \varphi_i(\bar{\theta}_i) \mathbb{1}_{i \neq j'})_{i \in N}$ and $z'_\epsilon := ([\varphi_i(p) - \epsilon] \mathbb{1}_{i=j} + \varphi_i(\bar{\theta}_i) \mathbb{1}_{i \neq j})_{i \in N}$. Observe that z_ϵ and z'_ϵ are outside of Z^* , but their midpoint is in Z^* when ϵ is sufficiently small—contradicting the convexity of $Z \setminus Z^*$.

Thus, $|\mathcal{J}| \leq 1$, and $|J| \leq 1$ for any $J \in \mathcal{J}$. It follows that \mathcal{J} is equal to \emptyset , $\{\emptyset\}$, or $\{\{i\}\}$ for some $i \in N$, delivering the proposition. \square

A.3.3. Proofs for Section 5.3

Proof of Proposition 5. Recall $N \in \mathbb{N}$ denotes both the number of agents and the set of agents. Because the set of optimal voting weight vectors is clearly convex and invariant to permutations, it follows that the uniform voting weight vector $\omega^N := \frac{1}{N} \mathbb{1}_N \in \Delta N$ is optimal. Hence, defining $\mathbf{x}^N := \mathbb{1}_{\frac{1}{N} \sum_{i \in N} \varphi_i \geq c}$, the seller's optimal profit is $\mathbb{E}[\mathbf{x}^N(\varphi_1 - c)]$. Note that $\mathbb{E}[\varphi_1] = \underline{\theta}_1$, and let $\zeta \in (0, \infty)$ be some uniform upper bound on the bounded function $|\varphi_1 - c|$.

To start, let us establish the result for the (easier) case in which $\underline{\theta}_1 \neq c$. Note that a unanimous posted price of $\bar{\theta}_1$ [resp. $\underline{\theta}_1$] attains a profit of 0 [resp. $\underline{\theta}_1 - c$]. Hence our the limiting value from an optimal collective posted price will necessarily coincide with the seller's limiting profit from an optimal mechanism, if we show that the latter attains the specified limit. Observe now that the law of large numbers implies $\mathbb{E}[\mathbf{x}^N]$ converges to 0 [resp. 1] as $N \rightarrow \infty$ if $\underline{\theta}_1 < c$ [resp. $\underline{\theta}_1 > c$]. Using this fact, let us show that the seller's optimal value converges to the desired quantity in each of these two cases. First, if $\underline{\theta}_1 < c$, then the distance between the seller's optimal profit and zero is

$$|\mathbb{E}[\mathbf{x}^N(\varphi_1 - c)]| \leq \zeta \mathbb{E}[\mathbf{x}^N] \rightarrow 0 \text{ as } N \rightarrow \infty.$$

Next, if $\underline{\theta}_1 > c$, then the distance between the seller's optimal profit and $\underline{\theta}_1 - c$ is

$$|(\underline{\theta}_1 - c) - \mathbb{E}[\mathbf{x}^N(\varphi_1 - c)]| = |\mathbb{E}[(1 - \mathbf{x}^N)(\varphi_1 - c)]| \leq \zeta \mathbb{E}[1 - \mathbf{x}^N] \rightarrow 0 \text{ as } N \rightarrow \infty,$$

as required.

We now turn to the complementary case that $c = \underline{\theta}_1$. As this equality implies $\varphi_1(\underline{\theta}_1) < c < \varphi_1(\bar{\theta}_1)$, let $\hat{\theta}_1 := \varphi_1^{-1}(c)$, the unique profit-maximizing posted price for a single agent with value distribution distributed according to CDF F_1 . Let $\pi := \max_{p \in \mathbb{R}_+} (p - c)[1 - F_1(p)] = (\hat{\theta}_1 - c)[1 - F_1(\hat{\theta}_1)]$, which is easily seen to be an upper bound on the seller's profit.²¹ Let us show, given an arbitrary $\epsilon > 0$, that our seller's optimal limit optimal profit, as $N \rightarrow \infty$, is at least $\pi - \epsilon$. Because F_1 is continuous, some $\theta_L, \theta_H \in \Theta_1$ exist such that $\theta_L < \hat{\theta}_1 < \theta_H$ and $F_1(\theta_H) - F_1(\theta_L) \leq \frac{\epsilon}{\zeta}$.

Given $N \in \mathbb{N}$ with $N > 1$, let X_1^N denote the interim allocation rule for agent 1 in the optimal mechanism. Observe now that

$$\begin{aligned} X_1^N(\theta_H) - X_1^N(\theta_L) &= \mathbb{P} \left[-\frac{1}{N} \varphi_1(\theta_H) \leq \frac{1}{N} \left(\sum_{i=2}^N \varphi_i \right) - c \leq -\frac{1}{N} \varphi_1(\theta_L) \right] \\ &\geq \mathbb{P} \left[\frac{-\delta}{N-1} \leq \left(\frac{1}{N-1} \sum_{i=2}^N \varphi_i \right) - c \leq \frac{\delta}{N-1} \right] \\ &\geq 1 - \gamma_N, \end{aligned}$$

for $\gamma_N := 2e^{\frac{-2\delta^2}{(N-1)(\theta_1 - \underline{\theta}_1)^2}}$, where the last inequality follows from Hoeffding's inequality. Interim monotonicity then tells us any $\theta_1 \in \Theta_1$ has $X_1^N(\theta_1) \leq \gamma_N$ if $\theta_1 \leq \theta_L$ and $X_1^N(\theta_1) \geq 1 - \gamma_N$ if $\theta_1 \geq \theta_H$. Therefore, the difference between π and the optimal profit is

$$\begin{aligned} \pi - \mathbb{E} [X_1^N(\theta_1)(\varphi_1 - c)] &= \mathbb{E} \{ [\mathbb{1}_{\theta_1 \geq \hat{\theta}_1} - X_1^N(\theta_1)] (\varphi_1 - c) \} \\ &\leq \mathbb{E} \{ [\gamma_N \mathbb{1}_{\theta_1 \notin (\theta_L, \theta_H)} + \mathbb{1}_{\theta_1 \in (\theta_L, \theta_H)}] |\varphi_1 - c| \} \\ &\leq \mathbb{E} [\gamma_N + \mathbb{1}_{\theta_1 \in (\theta_L, \theta_H)}] \zeta \\ &= [\gamma_N + F_1(\theta_H) - F_1(\theta_L)] \zeta \\ &\leq \zeta \gamma_N + \epsilon \\ &\rightarrow \epsilon \text{ as } N \rightarrow \infty. \end{aligned}$$

As $\epsilon > 0$ was arbitrary, it follows that the optimal profit approaches π as $N \rightarrow \infty$.

Now, let us see that the highest profit $\max_{p \geq 0} (p - c)[1 - F_1(p)]^N$ attainable with a collective posted price mechanism goes to 0 as $N \rightarrow \infty$. First, by Proposition 2, we can express this quantity as $\pi_N := \max_{p \in \Theta_1} (p - c)[1 - F_1(p)]^N$. Now, let $p_N \in [\underline{\theta}_1, \bar{\theta}_1]$ be a maximizer of $p \mapsto (p - c)[1 - F_1(p)]^N$, which exists by compactness and continuity. Because $[\underline{\theta}_1, \bar{\theta}_1]$ is compact, we can decompose the sequence $(p_N)_N$ into a collection of convergent subsequences $(p_{N_\ell})_\ell$. Consider an arbitrary such subsequence converging to p_* . The sequences $(p_{N_\ell} - c)_\ell$ and $([1 - F_1(p_{N_\ell})]^{N_\ell})_\ell$ are both bounded, and one of them converges to zero—the former if $p_* = \underline{\theta}_1$, the latter if $p_* > \underline{\theta}_1$. Hence the product converges to zero. \square

²¹Consider a relaxed program in which IC and IR are required only for one agent. The optimal profit in the relaxed program, a single-agent monopolist problem, is π .

A.4. Proofs for Section 6

LEMMA 5: For any $\gamma \in [0, 1]$ and $\lambda \in \Delta N$, an essentially unique allocation rule x^* solves

$$\begin{aligned} \max_{x \in \mathcal{X}} \min_{i \in N} \mathbb{E} \{x(\boldsymbol{\theta}) [(1 - \gamma)\boldsymbol{\varphi}_i + \gamma\lambda \cdot \boldsymbol{\theta} - c]\} & \quad (\text{PO}) \\ \text{s.t. } x \text{ is interim-monotone.} & \end{aligned}$$

It is given by $x^*(\boldsymbol{\theta}) := \mathbb{1}_{(1-\gamma)\omega \cdot \boldsymbol{\varphi} + \gamma\lambda \cdot \boldsymbol{\theta} \geq c}$, where $\omega \in \Delta N$ is any vector of weights such that $\text{supp}(\omega) \subseteq \text{argmin}_{i \in N} \mathbb{E} [\boldsymbol{\varphi}_i \mathbb{1}_{(1-\gamma)\omega \cdot \boldsymbol{\varphi} + \gamma\lambda \cdot \boldsymbol{\theta} \geq c}]$, and some such ω exists.

Proof. To establish this lemma, one can follow the proof of Theorem 1 with a modified objective, but otherwise verbatim. We look at a zero-sum game in which a Maximizer chooses from $\tilde{\mathcal{X}}$ (the set \mathcal{X} modulo almost-sure equivalence), a Minimizer chooses from ΔN , and the former's objective is

$$(x, \omega) \mapsto \mathbb{E} \{x(\boldsymbol{\theta}) [(1 - \gamma)\omega \cdot \boldsymbol{\varphi} + \gamma\lambda \cdot \boldsymbol{\theta} - c]\}.$$

Just as in the proof of Theorem 1: This game exhibits some Nash equilibrium by a compactness argument, and each Minimizer strategy $\omega \in \Delta N$ admits a unique best response $\mathbb{1}_{(1-\gamma)\omega \cdot \boldsymbol{\varphi} + \gamma\lambda \cdot \boldsymbol{\theta} \geq c}$ by the Maximizer. Hence, letting ω be some Nash equilibrium strategy for the Minimizer, it follows that $\mathbb{1}_{(1-\gamma)\omega \cdot \boldsymbol{\varphi} + \gamma\lambda \cdot \boldsymbol{\theta} \geq c}$ is the unique solution to the relaxation of program (PO) in which \mathcal{X}_{IM} is replaced with \mathcal{X} . As types are independent and (by regularity) $(1 - \gamma)\omega_i \varphi_i(\theta_i) + \gamma\lambda_i \theta_i$ is increasing in $\theta_i \in \Theta_i$, this solution is monotone and so solves (PO)—essentially uniquely so. Moreover, the condition satisfied by ω is exactly the Minimizer's best response property. The result follows. \square

Proof of Theorem 3. Let us begin with some payoff calculations. For each $\gamma \in [0, 1]$, $\lambda \in \Delta N$, and $i \in N$, define the function

$$\begin{aligned} g_i^{\gamma, \lambda} : \mathcal{X} & \rightarrow \mathbb{R} \\ x & \mapsto \mathbb{E} \{x(\boldsymbol{\theta}) [(1 - \gamma)\boldsymbol{\varphi}_i + \gamma\lambda \cdot \boldsymbol{\theta} - c]\}. \end{aligned}$$

If (x, m) is an IC mechanism generating a profile $\underline{U} \in \mathbb{R}^N$ of low-type utilities and a vector $v \in \mathbb{R}^{N+1} = \mathbb{R}^N \times \mathbb{R}$, then that the revenue can be computed equivalently as $\mathbb{E} [x(\boldsymbol{\theta})(\boldsymbol{\varphi}_j - c)] - \underline{U}_j$ for every $j \in N$ implies every $i \in N$ has

$$\begin{aligned} (\gamma\lambda, 1) \cdot v & = \mathbb{E} [x(\boldsymbol{\theta})(\boldsymbol{\varphi}_i - c)] - \underline{U}_i + \gamma\lambda \cdot \{\mathbb{E} [x(\boldsymbol{\theta})(\boldsymbol{\theta} - \boldsymbol{\varphi})] + \underline{U}\} \\ & = (1 - \gamma) \{\mathbb{E} [x(\boldsymbol{\theta})\boldsymbol{\varphi}_i] - \underline{U}_i\} + \gamma\lambda \cdot \{\mathbb{E} [x(\boldsymbol{\theta})\boldsymbol{\varphi}] - \underline{U}\} \\ & \quad + \gamma\lambda \cdot \{\mathbb{E} [x(\boldsymbol{\theta})(\boldsymbol{\theta} - \boldsymbol{\varphi})] + \underline{U}\} - c\mathbb{E} [x(\boldsymbol{\theta})] \\ & = g_i^{\gamma, \lambda}(x) - (1 - \gamma)\underline{U}_i. \end{aligned}$$

Moreover, as shown in Lemma 2, any interim-monotone allocation rule x can be paired with some transfer rule m for an IC and IR mechanism with low type utility $\underline{U}_{i_*} = 0$ for some $i_* \in N$ —which would then generate a payoff vector v with $(\gamma\lambda, 1) \cdot v = \min_{i \in N} g_i^{\gamma, \lambda}(x)$.

With these payoff calculations in hand, we now proceed to prove the theorem. Let \mathcal{X}_{IM} denote the set of interim-monotone allocation rules, which Lemma 2 points out are exactly those that can be used in an IC mechanism. In light of Lemma 5, it suffices to show that an allocation rule $x^* \in \mathcal{X}_{IM}$ is used in some Pareto optimal mechanism if and only if it belongs to $\operatorname{argmax}_{x \in \mathcal{X}_{IM}} \min_{i \in N} g_i^{\gamma, \lambda}(x)$ for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$.

First, suppose $x^* \in \operatorname{argmax}_{x \in \mathcal{X}_{IM}} \min_{i \in N} g_i^{\gamma, \lambda}(x)$ for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$. As noted in the above payoff calculations, some transfer rule m^* is such that the mechanism (x^*, m^*) is IC and IR and generates a payoff vector $v^* \in \mathbb{R}^{N+1}$ with $(\gamma\lambda, 1) \cdot v^* = \min_{i \in N} g_i^{\gamma, \lambda}(x^*)$. To show this mechanism is Pareto optimal, let (x, m) be any alternative IC and IR mechanism that generates a payoff vector $v \neq v^*$; our aim is to show that $v \not\geq v^*$. We proceed in two cases. First, if $x(\theta) = x^*(\theta)$ almost surely, then the revenue computations in the proof of Lemma 2 show $v_{N+1}^* - v_{N+1} = -(v_i^* - v_i)$ for every $i \in N$. Hence, in this case, we cannot have $v \geq v^*$. Second, suppose $x(\theta)$ and $x^*(\theta)$ are not almost surely equal. Then, letting $\underline{U} \in \mathbb{R}^N$ denote the profile of low-type interim utilities generated by (x, m) , we have

$$(\gamma\lambda, 1) \cdot v = \min_{i \in N} \left\{ g_i^{\gamma, \lambda}(x) - \underline{U}_i \right\} \leq \min_{i \in N} g_i^{\gamma, \lambda}(x) < \min_{i \in N} g_i^{\gamma, \lambda}(x^*) = (\gamma\lambda, 1) \cdot v^*,$$

where the strict inequality follows from the uniqueness property in Lemma 5. Because $(\gamma\lambda, 1) \geq 0$, it follows that $v \not\geq v^*$, as desired.

All that remains is to show every Pareto-optimal mechanism (x^*, m^*) solves program (PO) for some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$. To see this feature, observe the set $V \subseteq \mathbb{R}^{N+1}$ of attainable payoff vectors from IC and IR mechanisms is convex, because the set of IC and IR mechanisms is itself convex. Hence, the Minkowski sum $\hat{V} := V - \mathbb{R}_+^{N+1}$ is convex as well. Letting v^* denote the payoff vector generated by mechanism (x^*, m^*) , it follows that v^* is on the Pareto frontier of \hat{V} as well. In particular, v^* lies on the boundary of \hat{V} because $v^* + (\epsilon, \dots, \epsilon, \epsilon) \notin \hat{V}$ for any $\epsilon > 0$. Hence, the supporting hyperplane theorem delivers some nonzero $\tilde{\lambda} \in \mathbb{R}^{N+1}$ such that $v^* \in \operatorname{argmax}_{v \in \hat{V}} \tilde{\lambda} \cdot v$. Observe, every $i \in N$ has $\tilde{\lambda}_i \geq 0$, for otherwise $v = v^* - e_i \in \hat{V}$ would have $\tilde{\lambda} \cdot v > \tilde{\lambda} \cdot v^*$; and $\tilde{\lambda}_{N+1} \geq \sum_{i \in N} \tilde{\lambda}_i$, for otherwise $v = v^* + (1, \dots, 1, -1) \in \hat{V}$ (which is attainable by lowering the transfer rule by constant 1) would have $\tilde{\lambda} \cdot v > \tilde{\lambda} \cdot v^*$. Together, that $\tilde{\lambda}_i \geq 0$ for every $i \in N$ and $\tilde{\lambda}_{N+1} \geq \sum_{j \in N} \tilde{\lambda}_j$ tell us $\tilde{\lambda}_{N+1} > 0$ because $\tilde{\lambda} \neq 0$. Hence, rescaling $\tilde{\lambda}$ if necessary, we may assume $\tilde{\lambda}_{N+1} = 1$. Then, because $\tilde{\lambda}_i \geq 0$ for every $i \in N$ and $\sum_{j \in N} \tilde{\lambda}_j \leq 1$, it follows that some $\gamma \in [0, 1]$ and $\lambda \in \Delta N$ exist for which $\tilde{\lambda} = (\gamma\lambda, 1)$.

Let us now show x^* solves (PO), for the given γ and λ . Assuming otherwise, for a contradiction, let $x \in \mathcal{X}_{IM}$ attain a strictly higher objective in program (PO). As noted above, some transfer rule m exists such that (x, m) is IC and IR with low type utility $\underline{U}_{i_*} = 0$ for some $i_* \in N$. Hence, letting \underline{U}^* denote the profile of low-type utilities generated by (x^*, m^*) , and $v \in V$ denote the payoff vector

generated by mechanism (x, m) , we have

$$\tilde{\lambda} \cdot v = g_{i_*}^{\gamma, \lambda}(x) \geq \min_{i \in N} g_i^{\gamma, \lambda}(x) > \min_{i \in N} g_i^{\gamma, \lambda}(x^*) \geq \min_{i \in N} \left\{ g_i^{\gamma, \lambda}(x^*) - (1 - \gamma) \underline{U}_i^* \right\} = \tilde{\lambda} \cdot v^*,$$

a contradiction. Therefore, x^* solves program **(PO)**. □