Delegated Contracting*

João Thereze† Udayan Vaidya‡

January 15, 2026

Abstract

A principal seeks to contract with an agent but must do so through an informed delegate. Although the principal cannot directly mediate the interaction, she can constrain the menus of contracts the delegate may offer. We show that the principal can implement any outcome that is implementable through a direct mechanism satisfying dominant strategy incentive compatibility and ex-post participation for the agent. We apply this result to several settings. First, we show that a government that delegates procurement to a budget-indulgent agency should delegate an interval of screening contracts. Second, we show that a seller can delegate sales to an intermediary without revenue loss, provided she can commit to a return policy. Third, in contrast to centralized mechanism design, we demonstrate that no partnership can be efficiently dissolved in the absence of a mediator. Finally, we discuss when delegated contracting obstructs efficiency, and when choosing the right delegate may help restore it.

*We thank Leslie Marx, Andrew McClellan, and Harry Pei for useful comments and suggestions. We also thank the seminar audiences at Arizona State, Duke, Stanford, NC State, UBC, UNC and UC Riverside for helpful feedback.

†The Fuqua School of Business, Duke University, Durham, NC 27708 — joao.thereze@duke.edu.

‡The Fuqua School of Business, Duke University, Durham, NC 27708 — udayan.vaidya@duke.edu.

1 Introduction

In many economic relationships, a principal must rely on an intermediary to contract with an agent on her behalf. Although the principal does not interact directly with the agent, she may retain the authority to limit the intermediary's contractual discretion. Governments, for example, routinely delegate procurement decisions to agencies, subject to regulations that determine which procurement contracts are admissible. Collectors consign their artworks to auction houses, who sell to final buyers using agreed-upon sales formats. Manufacturers impose resale price maintenance on retailers, enforcing floors or ceilings on prices charged to consumers. We collectively refer to these environments as instances of *delegated contracting*.

The prevalence of delegated contracting raises a basic problem of institutional design: how can a principal shape outcomes when authority must be exercised indirectly? When communication with the downstream party is infeasible, the principal must rely on *ex-ante* restrictions over the intermediary's contracting process. While the principal could give the intermediary no flexibility, doing so would prevent the intermediary from exploiting his superior information. How should the allocation of contractual rights balance the principal's control with the intermediary's flexibility? What is the cost of this loss of control for the principal? What are the distortions induced by contractual delegation? These questions are central to the design of organizations, supply chains, and regulation.

We study how to delegate authority in settings where centralized implementation is infeasible. Our contribution is twofold. At a conceptual level, we provide a general framework that identifies the principal's loss from delegating the contractual interaction, rather than centralizing it. At a practical level, we derive the principal's optimal contractual restrictions in several economically relevant settings.

Model and Main Theorem A principal seeks to contract with an agent but must do so through a delegate. Although both the delegate and the agent hold private information relevant to the interaction, the principal cannot communicate directly with the agent. The principal's only instrument is to impose ex-ante restrictions on the menus of contracts the delegate may offer. We assume menus are *transparent* (Zheng, 2002): given the menu, the agent's choice of contract alone determines the outcome. After observing the constraints imposed by the principal, the delegate selects an admissible menu of contracts and presents it to the agent, who chooses whether to sign a contract or walk away. This choice determines the payoffs of all three parties. Our baseline model assumes a private values environment, though our results extend to more general preferences and information structures.[1]

Our main result, Theorem 1, characterizes which social choice rules—mappings from the parties' private information to outcomes—can arise through contractual delegation. We show that, through delegation, the principal can implement any social choice rule that is implementable in a centralized Bayesian mechanism, subject to dominant-strategy incentive compatibility and ex-post individual rationality for the agent. The stronger incentive constraints on the agent's side reflect the sequentiality of the delegated environment: the delegate's choice of menu may signal private information that the agent can respond to. Under transparency and private values, there is no loss in the delegate's menu fully revealing his private information. As a result, agent incentive compatibility must hold uniformly across the delegate's types.

Our characterization identifies the principal's loss from delegating contractual authority. Delegation restricts the principal to a narrower set of social choice rules compared to centralized mechanism design. The gap between dominant and Bayesian incentives for the agent captures the cost of delegation. Beyond its conceptual implications, this characterization yields a practical benefit: it recasts the design problem under delegation as a more familiar mechanism design exercise, subject to tighter constraints on the agent. This formulation opens the door to applying existing tools and insights from the broader literature to solve for optimal contractual restrictions in delegated contracting environments.

Applications We illustrate our framework through several applications. In each case, we use Theorem 1 to identify the principal-optimal mechanism and characterize the contractual restrictions that support its delegated implementation. These examples demonstrate how our general setup can be applied to concrete institutional settings and clarify when the cost of delegation is negligible and when it is substantial.

Our first application studies the problem of a legislature delegating procurement of goods or services to a budget-indulgent agency. The agency privately observes the benefits of procurement and contracts with a supplier whose cost is unknown. The legislature must design contractual restrictions that simultaneously screen the agency and allow the agency to screen the supplier. Under standard regularity assumptions, the legislature's optimal policy delegates an *interval of contracts*, one for each type of the agency up to a cap. These contracts coincide with the agency's optimal screening contracts un-

[1] See Section 2.2 for a discussion on how our main result generalizes to certain common value environments.

der full flexibility, so that agencies whose benefit from procurement does not exceed the cap are unrestricted. Implementation is straightforward: the agency may offer any menu composed of contracts lying below a prescribed price-quality frontier. Moreover, the optimality of interval delegation readily generalizes to more complex procurement problems involving multiple suppliers, multi-dimensional supplier private information, or richer screening mechanisms available to the agency.

We next examine a revenue-maximizing seller who must reach a final buyer through a privately informed intermediary. This intermediary can resell the product or use it for private benefit. We show that the seller can achieve the same revenue through delegation as in an optimal auction involving the intermediary and the buyer. This optimal mechanism can be implemented through a resale price agreement and buyback policy. By tying the intermediary's resale price to the seller's "wholesale" price, the seller induces the intermediary to transfer the good to the buyer whenever the seller finds it optimal. The buyback policy encourages intermediation: without it, intermediaries with low private valuations for the good may refuse to purchase in the first place. When buybacks are infeasible, the seller's optimal mechanism involves only a price agreement, which is discounted to make sure that all intermediaries purchase. Thus, the seller's desire to reach downstream buyers dominates her incentive to exercise monopoly power, and delegated contracting improves the seller and buyer surplus relative to a laissez-faire benchmark.

Third, we ask whether delegation obstructs the implementation of efficient outcomes. We consider a standard quasilinear environment in which outside options are normalized to zero. When the players' types are substitutes from the perspective of welfare—that is, in a rival allocation problem—delegated contracting implements the efficient allocation subject to budget balance. The contracts resemble a pivot mechanism, modified to ensure budget balance and dominant strategies for the agent. In contrast, in non-rival allocation problems, the efficient allocation cannot be implemented with budget balance. These results parallel centralized mechanism design (Krishna and Perry, 1998; Mailath and Postlewaite, 1990), illustrating a class of problems in which there is no loss from contractual delegation.

Our final application revisits the classic partnership dissolution problem, in which a principal seeks to efficiently reallocate ownership of a jointly owned asset subject to budget balance (Cramton et al., 1987). The partners' outside options, as determined by their ownership shares, are critical: in a centralized mechanism, efficient trade is feasible whenever the initial ownership shares are not too unequal. Under delegated contracting, however, efficiency is unattainable regardless of the initial ownership distribution. The reason for this disparity is that the delegate is not willing to sell his share of the asset efficiently. On one hand, efficiency means that a low-valuation delegate has to part with his holdings almost surely; on the other, ex-post individual rationality for the agent means that the price at which this delegate must sell his shares is quite low. Facing sure expropriation with almost no compensation, the delegate prefers to walk away. This sharp contrast highlights the limits of delegated contracting and the value of mediation in centralized design.

In this context, we then consider how the identity of the delegate affects outcomes, i.e., *whom to delegate to*. Introducing losses from diffuse ownership—where the payoff from a partial stake is less than proportional to full ownership—lowers the outside options and allows efficiency to be achievable, depending on which partner holds delegated authority. A natural conjecture is that the principal should delegate to the player with a larger ownership share, as the delegate's interim participation constraint is easier to satisfy than the agent's ex-post constraint. We show, however, that the reverse holds: whenever delegating to one player is dominant, that player holds a smaller share of the asset. This insight runs counter to conventional wisdom, suggesting control rights should be inversely related to ownership stakes, provided the appropriate contractual restrictions.

Related Literature Our paper contributes to the literature on hierarchical contracting stemming from Tirole (1986), which examines the distortions that arise when a principal anticipates subcontracting among agents.[2] These papers study the cost of delegation when the principal can observe individual outputs (Mookherjee, 2012), restrict the outcome space (Malladi, 2022), and design the players' outside options (Dworczak and Muir, 2024). Relative to this literature, our principal is able to exercise tighter control on the downstream interaction, as the subcontracting arrangement is itself contractible. Closest to us, Bhaskar et al. (2023) study a regulator who restricts the menus an insurer may offer to a consumer. They show that latent contracts allow the regulator to freely extract the information commonly known to both the insurer and consumer. We can reinterpret their result within our framework: because shared information can be elicited at no cost to the principal in a centralized Bayesian mechanism, latent contracts

[2] One focus of this literature is anticipating "collusion" between the agents due to unobserved subcontracting (Baliga and Sjöström, 1998; Laffont and Rochet, 1997; McAfee and McMillan, 1995; Mookherjee and Tsumagari, 2004; Severinov, 2003; Tirole, 1992). See Mookherjee (2006) and Mookherjee (2012) for reviews of this literature. Because our principal observes and can contract on downstream contracting, she is not concerned with collusion *per se*. Yet, delegation remains costly due to the principal's inability to communicate directly with the downstream agent.

serve to implement this centralized benchmark in dominant strategies for the consumer.[3]

Second, we connect to the mechanism design literature that links centralized and decentralized implementations. Several papers identify settings in which a delegated implementation replicates the optimal centralized mechanism (Bhaskar et al., 2023; Faure-Grimaud and Martimort, 2001; Melumad et al., 1992; Melumad and Shibano, 1991; Melumad et al., 1995; Mitchell, 2025). Our main result tightens this connection between delegated contracting and centralized mechanism design. We identify exactly when the optimal centralized mechanism can be implemented in a delegated fashion. Further, Theorem 1 provides the program to solve for the optimal contractual rights when the centralized benchmark is unattainable, quantifying the principal's loss from delegation.

Finally, we build on the delegation literature, which explores how a principal should allocate decision rights to a informed but biased expert (Alonso and Matouschek, 2008; Holmström, 1980). In the canonical model, the principal chooses how much discretion to grant an expert who chooses an action from an interval. This work shows that interval delegation—restricting the delegate choice to a censored interval of actions—often suffices when the action space is one-dimensional and the expert's bias is state-independent (Alonso and Matouschek, 2008; Amador and Bagwell, 2013, 2022; Guo, 2016; Hiriart and Martimort, 2012; Kolotilin and Zapechelnyuk, 2025). However, in richer environments where decisions and information are multidimensional, the optimal delegation set may take more complex forms (Ambrus and Egorov, 2017; Frankel, 2014; Koessler and Martimort, 2012). One of our contributions is to extend the interval delegation logic to settings where the delegate chooses not a single action but a menu of contracts offered to a downstream agent. In particular, we show that in some environments—such as procurement—the notion of interval delegation generalizes naturally to the space of menus. Closer to our exercise, Kundu and Nilssen (2020) and Martimort et al. (2020) consider models where the delegate interacts with downstream agents, but restrict the principal to a simple class of constrained delegation mechanisms. We make no such restriction.

2 A model of delegated contracting

A principal (she) seeks to achieve an outcome through the interaction between two privately informed players, indexed by $i \in \{1,2\}$. We refer to player 1 as the *delegate* and player 2 as the *agent*. Each player i holds private information or type θ_i, with (θ_1,θ_2)

[3]Though our baseline model considers a private values setting, Theorem 1 holds verbatim under one-directional common values where the agent's type is payoff-relevant for the delegate, but not vice versa.

drawn jointly from $\Theta := \Theta_1 \times \Theta_2$ according to a probability measure $\mu_o \in \Delta\Theta$. The set of outcomes is denoted by X, and player i's utility from outcome $x \in X$ given private information θ_i is $u_i(x, \theta_i) \in \mathbb{R}$. If no agreement is reached, a default outcome $o \in X$ occurs.

The principal cannot directly observe the interaction between the delegate and the agent. Instead, she entrusts contractual authority to the delegate, who interacts privately with the agent. Yet, the principal retains control over the contract space by specifying the allowable menus of contracts. Formally, a *contract* is simply an outcome $x \in X$ agreed upon by the delegate and the agent, e.g., a price-quantity pair. The timing of the game is as follows.

1. The principal specifies the delegate's contractual rights, $\mathcal{C} \subseteq 2^X$. These contractual rights are a collection of subsets of X, representing the admissible menus of contracts that can be offered to the agent.

2. The delegate observes his type θ_1 and either selects an allowable menu $C \in \mathcal{C}$ to offer or opts out, yielding outcome o.

3. The agent observes θ_2 and chooses to sign an available contract $x \in C$ or to opt out, yielding outcome o.

We now define the strategies in the subgame between the delegate and agent, given the principal's choice of contractual rights $\mathcal{C}$. A strategy for the delegate maps the delegate's information to a lottery over menus, $\sigma_1 : \Theta_1 \to \Delta(\mathcal{C} \cup \{\{o\}\})$. A strategy for the agent, given the delegate's menu C, maps the agent's information to a lottery over contracts, $\sigma_2^C : \Theta_2 \to \Delta(C \cup \{o\})$. The agent's beliefs over the delegate's type after observing the menu C are given by $\mu^C \in \Delta\Theta_1$.

A *Perfect Bayesian Equilibrium* (henceforth, an equilibrium) given contractual rights $\mathcal{C}$ is a tuple $\left(\sigma_1, \{\sigma_2^C, \mu^C\}_{C \in \mathcal{C}}\right)$ such that:

(i) $\sigma_2^C(\theta_2)$ maximizes the agent's expected payoff given belief μ^C, for every menu C and type θ_2;

(ii) $\sigma_1(\theta_1)$ maximizes the delegate's expected payoff given σ_2^C, for every type θ_1;

(iii) Beliefs μ^C are derived from the delegate's strategy using Bayes' rule whenever possible.

2.1 Delegated Implementation

Our first goal is to establish what outcomes can be obtained through delegated contracting. Formally, a social choice function is a mapping between the private information of both players and the set of outcomes, $Y : \Theta \to X$.[4] A social choice function is *implementable through delegation* if there exist contractual rights $\mathcal{C}$ and an associated equilibrium $\left(\sigma_1, \{\sigma_2^C, \mu^C\}_{C\in\mathcal{C}}\right)$ leading to outcomes given by Y. That is, given the strategies used by the delegate and the agent with types $(\theta_1,\theta_2) \in \operatorname{supp}\mu_o$, outcome $Y(\theta_1,\theta_2)$ arises with probability 1:

$$\int_{C\in\mathcal{C}} \left[\sigma_2^C(\theta_2)\left(Y(\theta_1,\theta_2)\right)\right]\sigma_1(\theta_1)(dC) = 1.$$

We compare the outcomes implementable through delegation with those achievable via *centralized mechanisms*, in which the principal directly communicates with both parties. A centralized mechanism is a pair (M,g) with $M = M_1 \times M_2$ and $g : M \to \Delta X$, where M is a message space and g is an outcome function that maps messages into a chosen contract. The mechanism (M,g) is the direct mechanism associated with Y if $M = \Theta$, and $g = Y$.

We start by restating standard definitions in the mechanism design literature. First, a social choice function is said to be Bayesian incentive compatible (BIC) and interim individually rational (IIR) if, for $i \in \{1,2\}$, and all $\theta_i, \theta' \in \Theta_i$:

$$\mathbb{E}_{\theta_{-i}}\left[u_i\left(Y(\theta_i,\theta_{-i}),\theta_i\right)|\theta_i\right] \geq \max\left\{\mathbb{E}_{\theta_{-i}}\left[u_i\left(Y(\theta',\theta_{-i}),\theta_i\right)|\theta_i\right], u(o,\theta_i)\right\}.$$

As usual, a social choice function is BIC if it's corresponding direct mechanism induces truthful reporting, given each player's beliefs about the other's type. It is IIR if both players prefer to participate in this mechanism than to obtain the disagreement outcome. A social choice function is said to be dominant-strategy incentive compatible (DSIC) and ex-post individually rational (EPIR) for party $i \in \{1,2\}$ if for all $\theta_i, \theta \in \Theta_i$, and $\theta_{-i} \in \Theta_{-i}$:

$$u_i\left(Y(\theta_i,\theta_{-i}),\theta_i\right) \geq \max\left\{u_i\left(Y(\theta',\theta_{-i}),\theta_i\right), u(o,\theta_i)\right\}.$$

It is well-known that, by the revelation principle, a social choice function Y is implementable in a centralized mechanism if it satisfies BIC and IIR for both the delegate and the agent (Myerson, 1981). Our main result characterizes how contractual delegation

[4] The assumption that the social choice function is deterministic is without loss of generality in our model. Concretely, we impose no restriction in the outcome space, so one can assume that it includes lotteries and our results are unchanged. We discuss this assumption in the Discussion section.

narrows the set of outcomes that can be achieved compared to centralized mechanism design. All the proofs are in the Appendix.

Theorem 1. *A social choice function Y is implementable through delegation if and only if it is implementable in a centralized mechanism and satisfies:*

(i) Dominant-strategy incentive compatibility (DSIC) for the agent;

(ii) Ex-post individual rationality (EPIR) for the agent.

Theorem 1 shows that the principal can implement any outcome achievable under centralization, subject to one caveat: the agent must be willing to participate and report truthfully even after learning the delegate's type. A simple intuition driving this result is that the delegate's choice of menu signals his private information. Even though this information is payoff-irrelevant due to the transparency and private values assumptions, it allows the agent to respond differently to different delegate types. Thus, Bayesian incentives are insufficient. In fact, there is no loss in the delegate's menu fully revealing his type, requiring agent incentives type-by-type of the delegate.

This result is useful for two reasons. First, it provides a practical tool to study delegated contracting problems. Instead of solving the interaction between the delegate and the agent and proceeding by backward induction to search over the unstructured space of contractual rights, Theorem 1 translates delegated contracting into a set of familiar implementation constraints. Second, the result precisely delineates the loss from delegation compared to the standard, centralized mechanism design. In several applications, this characterization makes it relatively straightforward to quantify the loss from delegation. In doing so, Theorem 1 provides another tool to the literature, which identifies some environments featuring no loss from delegated contracting, but is largely silent about the extent of such losses when they arise.

2.2 Discussion of Assumptions

Menus of Contracts vs. Mechanisms A central assumption is that the delegate can only offer a menu of outcomes, rather than a full-fledged mechanism. The menus of outcomes can be equivalently interpreted as full mechanisms subject to a simple restriction: they may respond to the agent's message but not to the delegate's. In other words, the delegate cannot affect the outcome once he offers the menu. Such mechanisms in which a privately informed designer cannot affect the outcome have been called *transparent* in the literature (Zheng, 2002).

Transparency is appealing for two reasons. First, given the principal's limited ability to monitor the delegate, transparency ensures that the agent himself can audit the delegate. Because the agent can verify the mapping from message to outcome, this mitigates concerns over manipulation. Moreover, the transparency assumption guarantees tractability by avoiding the complexities of informed principal problems (Maskin and Tirole, 1990; Myerson and Satterthwaite, 1983). While in such models the principal might be able to expand the set of implementable outcomes by exploiting the agent's strategic uncertainty, they are considerably harder to analyze. Moreover, the ubiquity of menu-like contracting processes in practice suggests that transparent mechanisms are empirically relevant.

Private Values The private values assumption is meaningful because it rules out the delegate's signalling through menu choice: the delegate's selection of menu conveys no payoff-relevant information to the agent. Nonetheless, Theorem 1 still holds under common values as long as the agent's information about common-payoff-relevant states is a refinement of the delegate's information about those states. Formally, let $\theta_i = (\theta_i^p, \theta_i^c)$, and assume the utilities have the form $u_i(\theta_i^p, \theta_1^c, \theta_2^c)$. That is, θ_i^p represents the component of information that affects only the utility of type i, whereas θ_i^c affects both parties. If θ_2 is a sufficient statistic for θ_1^c, Theorem 1 holds verbatim. This is important because common-values settings in which the delegate is less informed than the agent are ubiquitous both in practice and in the literature (Bhaskar et al., 2023), and our model encompasses those environments.

Deterministic Social Choice Functions We assumed social choice functions are deterministic, but this is without loss of generality. There are two, equivalent, ways of accommodating lotteries over outcomes X in the co-domain of the social choice function. First, one could re-define the outcome space to be the space of lotteries over X and, subject to measurability assumptions, Theorem 1 applies directly. Alternatively, one could maintain the outcome space X, but redefine contracts as belonging to ΔX. In this case, the allowable menus consist of sets of lotteries, $\mathcal{C} \subseteq 2^{\Delta X}$. Embedding randomization into the contracts this way lets Theorem 1 apply to random social choice functions.

3 Applications

3.1 Delegated Procurement

We apply our delegated contracting framework to study a stylized procurement environment. In this setting, a principal delegates the task of contracting with a supplier (the agent) to an internal division (the delegate). For example, consider a federal agency or multinational corporation delegating purchase of equipment or services to regional offices. Delegated contracting entails a familiar tradeoff: the delegate has better knowledge of local conditions, but may be overindulgent with the principal's resources, leading to inefficient or excessive spending. How can the principal indirectly induce efficient procurement outcomes without directly intervening in the downstream negotiations?

The procurement task involves securing a good or service characterized by a one-dimensional quality or quantity level $q \in [0,1]$ delivered by the seller. A contract then specifies the associated transfer or payment $t \in \mathbb{R}_+$ made to the seller for any chosen quality level q. As in our general model, the game begins with the principal restricting the menus of contracts the delegate may offer. Afterwards, the delegate privately observes a benefit parameter $b \in [0,1]$, representing the per-unit value of procurement. Similarly, the agent's cost type is privately observed, parameterized by $s \in [0,1]$. We assume that $b \sim F$ and $s \sim G$, both of which admit strictly positive densities on $[0,1]$. We additionally assume that F satisfies the monotone hazard rate assumption that $\frac{1-F(b)}{f(b)}$ is non-increasing.

When purchasing q units at a total price of t, the payoffs are as follows. The principal's payoff is $U_P = bq - (1+\alpha)t$, where $\alpha \geq 0$ captures the degree of misalignment between the principal and delegate. The delegate's payoff is $U_D = bq - t$, where we assume the delegate benefits from procurement proportionally to b but does not fully internalize the cost to the principal. Thus, α reflects internal transfer frictions, budgetary costs, or other inefficiencies within the organization. The agent's payoff is $U_A = t - sc(q)$, where c is strictly increasing, strictly convex, and satisfies the Inada conditions that $c(0) = 0$, $c'(0) = 0$, $\lim_{q \to \infty} c(q) = \infty$.

Our focus is to identify the structure of the principal's optimally-delegated set of contracts. Since the principal cannot directly pay the agent, the delegation set must fulfill the dual role of preventing the delegate's opportunism and screening the agent's private cost. There are many potential tools to manage these conflicts, including price and quality caps, "taxing" the delegate to deter overspending, or bounds on the price-to-quality ratio. Applying Theorem 1 to this setting gives us the following description of the principal's problem.

$$
\begin{aligned}
\max_{q(b,s),t(b,s)} \quad & \mathbb{E}_{b,s}\left[bq(b,s) - t(b,s)\right] \\
s.t. \quad & \mathbb{E}_s\left[bq(b,s) - \alpha t(b,s)\right] \geq \mathbb{E}_s\left[bq(b',s) - \alpha t(b',s)\right] \quad \forall b,b' && \text{(BIC -Delegate)} \\
& t(b,s) - sc(q(b,s)) \geq t(b,s') - sc(q(b,s')) \quad \forall b,s,s' && \text{(DSIC - Agent)} \\
& \mathbb{E}_s\left[bq(b,s) - \alpha t(b,s)\right] \geq 0 \quad \forall b && \text{(IIR - Agent)} \\
& t(b,s) - sc(q(b,s)) \geq 0 \quad \forall b,s && \text{(EPIR - Seller)}
\end{aligned}
$$

To solve this problem, it turns out that it is useful to frame the principal's optimal menu of contracts in terms of the delegate's *optimal screening contracts*. The $b-$**optimal screening contract** of the type b delegate, which we denote by (q_b^*, t_b^*), maximizes the delegate's payoff when her type is b, subject to incentive compatibility and participation of the agent. In other words, this is the contract that the type-b delegate would choose if given full flexibility by the principal. Formally, the b-optimal contract is defined as the solution to:

$$
\begin{aligned}
\max_{q(s),t(s)} \quad & \mathbb{E}_s[bq(s) - t(s)] \\
s.t. \quad & t(s) - sc(q(s)) \geq t(s') - sc(q(s')) \\
& t(s) - sc(q(s)) \geq 0
\end{aligned}
$$

Proposition 1. *There exists a cutoff $\hat{b} < 1$ such that the principal delegates an interval of b-optimal screening contracts $(q_b^*, t_b^*)_{b \in [0,\hat{b}]}$. The principal's optimal menu can be implemented by allowing the delegate to choose any contract below the frontier given by $(q_{\hat{b}}^*, t_{\hat{b}}^*)$.*

This characterization delineates how the principal separately provides incentives to the delegate versus the downstream agent. Importantly, the principal optimally does not distort the offered contracts themselves, only the *set* of available contracts. All delegate types below $\hat{b}$ choose their optimal screening contract as if they were unconstrained, while those above $\hat{b}$ all choose the maximal contract $(q_{\hat{b}}^*, t_{\hat{b}}^*)$. Remarkably, this implies that even though the principal can restrict the entire contract that the delegate offers to the agent, it is sufficient to only restrict the set of feasible ex-post outcomes.

One (incomplete) intuition behind the structure of the optimal delegation set is that the allowable contracts must satisfy a particular cost-minimization property. Given that the principal and delegate both have preferences over the expected quality, when an agent procures an expected $\bar{Q}$ units, it is preferable for the delegate to do so in the cheapest way possible. Doing so can only relax the budget balance and IC constraints, as any additional slack can always be made up for by a lump sum transfer. This cost-minimization problem can be written as

$$\begin{aligned} &\min_{q,t} \mathbb{E}_s[t(s)] \\ s.t.\quad &\text{seller IC} \\ &\text{seller IR} \\ \mathbb{E}_s q(s) &= \bar{Q} \end{aligned}$$

A solution to this problem must also solve the following, where λ is the (endogenous) multiplier on the quality constraint:

$$\begin{aligned} &\max_{q,t} \mathbb{E}_s[\lambda q(s) - t(s)] \\ s.t.\quad &\text{seller IC} \\ &\text{seller IR} \end{aligned}$$

However, notice that the above problem is exactly the optimal screening problem for the type-λ delegate. Therefore, it is without loss for the principal to restrict attention to contracts that use efficient quality schedules.

One advantage of framing our discussion in terms of the delegate's optimal screening contracts is that the result fully generalizes regardless of how one parametrizes the downstream procurement problem. The agent's type need not be ordered in some single-crossing fashion, and the result also accommodates the possibility of multiple agents who compete in an auction. In any case, the principal should delegate an interval of optimal screening contracts (mechanisms). Of course, the specification of the downstream problem is needed to determine the b-optimal screening contracts themselves and the threshold type $\hat{b}$, but the underlying structure of the problem is unchanged. This generalization is facilitated by a particular relaxation of the principal's design problem which we use in the proof of the result. At a high level, we can abstract away the details of the downstream procurement problem by replacing them with an upper bound on the utility of the type b delegate: the utility delivered by the b-optimal screening contract. This relaxation is akin to a familiar one-dimensional delegation problem in which the principal would cap the agent's action, and we show that the relaxed optimum is also feasible in the original problem. The proof leverages results from Kleiner et al. (2021), as a majorization-like constraint naturally appears in the relaxed problem.

3.2 Delegated Contracts for Resale

We now study a stylized model of market intermediation. A seller (the principal) wishes to place a product in the market, but must do so through an informed intermediary (the delegate) who ultimately resells the product to a final consumer (the agent). In addition to the final consumer valuing the good, the good is valuable to the intermediary. If the intermediary retains the product, he might privately enjoy it, resell it later in a non-contractible manner, or scrap it for its "salvage value". The intermediary's private value is $\theta_1 \sim G \in \Delta[0,1]$, and the consumer's private value is $\theta_2 \sim F \in \Delta[0,1]$. To capture stochastic gains from trade between the intermediary and the final consumer, we assume that, for any price p, the consumers who are willing to pay that price have stronger valuation than the intermediaries willing to keep the product at that price. Formally, the distribution of consumer valuations conditional on any price stochastically dominates that of the intermediary: $F(v|v \geq p) \leq G(v|v \geq p)$, for all $p \in [0,1]$ and $v \geq p$. Finally, we assume F and G are continuous and have full support, with densities f and g respectively, and with an increasing hazard rate. The virtual value for distribution H with density h is defined as $\psi_H(x) = x - \frac{1-H(x)}{h(x)}$ and, by the assumption of increasing hazard rates, ψ_H is strictly increasing for $H \in \{F, G\}$. Both the intermediary and the consumer have quasilinear utilities with respect to money.

The timeline of the problem is as follows: first, the seller contracts with the intermediary, deciding whether to sell the product to the intermediary and constraining how it can be resold to the final consumer. Importantly, even if the intermediary buys the good initially, the seller might be willing to buy it back in case the intermediary fails to resell it to the final consumer. In the language of our model, the seller picks a contract space, $\mathcal{C} = \{C^\lambda\}_{\lambda \in \Lambda}$, in which each outcome encompasses the final allocation of the good between seller, intermediary and final consumer, and transfers to all the parties. At this point, we will impose no further restrictions on the set of contracts. Later in this section, we will investigate the effects of restrictions in the order of payments and on the buyback policy. We start by discussing the *laissez-faire* benchmark, in which the contract between the seller and the intermediary does not specify a resale policy. In this case, it is well-known that the intermediary with type θ_1 optimally sells the good by posting a monopoly price, $p_m(\theta_1)$, which satisfies:

$$p_m(\theta_1) = \psi_F^{-1}(\theta_1).$$

The intermediary's resulting profits, inclusive of his value when retaining the good are

$$\pi_m(\theta_1) = \psi_F^{-1}(\theta_1)\left(1 - F\left(\psi_F^{-1}(\theta_1)\right)\right) + \theta_1 F\left(\psi_F^{-1}(\theta_1)\right).$$

Therefore, from the point-of-view of the seller, it is as if she is selling a good to a population with valuation $v = \pi_m(\theta_1)$ distributed according to $H = G \circ \pi_m^{-1}$, which can be proved to also have strictly increasing virtual values. Thus, the seller sells to the intermediary at a *laissez-faire* posted price p_{LF} that is optimal against the distribution H.

Generically, *laissez-faire* contracting leads to two inefficiencies. First, the intermediary leverages his bargaining power to the extreme, acting as a full monopolist and retaining the good even when the valuation of the final consumer exceeds his own—that is, if $\theta_1 < \theta_2 \leq \psi_F^{-1}(\theta_1)$. This is an intensive margin distortion: when intermediaries hold the good, they post too high a price, foregoing some efficient trades. Second, the seller typically posts a price that excludes intermediaries with a lower valuation from the market. This is an extensive margin distortion, because some intermediaries are outright excluded from trading. Because this extensive margin distortion breaks the intermediation chain, the good may fail to reach high willingness-to-pay consumers. We will show that the optimal delegated contract space deals with both of these inefficiencies to some extent.

To determine how the seller should constrain the intermediary's resale, we appeal to Theorem 1. Although the physical description of our environment—in which the intermediary holds the good before the consumer—suggests that the seller might cede extra rents to the intermediary, Theorem 1 implies this is not the case. From the perspective of a centralized mechanism, the seller is an auctioneer facing two possible consumers. We know that the optimal auction allocates the good to the agent with the highest virtual value, when positive; otherwise, the seller retains the good. This asymmetric auction maximizes the seller's revenue under BIC and IIR constraints but, importantly, it is also DSIC and EPIR for both bidders. Theorem 1 then implies that the same revenue can be achieved by contractual delegation.

We now describe an indirect contract space that implements the optimal asymmetric auction as an equilibrium. Each allowable outcome can be thought of as a combination of two types of clauses: price agreements and buyback policies. A **price agreement** with resale price p is a procedure whereby the intermediary agrees with the seller to resell the good to the consumer at price p, and buys the good from the seller at some pre-specified discounted price $p - d(p)$. A **buyback policy** with refund r is a commitment from the seller to rebuy any unsold units at price r.

Proposition 2. *A set of menus $\mathcal{C}$ with the following two characteristics maximizes the seller's revenue.*

1. *Each menu contains a buyback policy with refund $r := \psi_G^{-1}(0)$.*
2. *Each menu $C^p \in \mathcal{C}$, $p \geq \underline{p} := \psi_F^{-1}(0)$, contains a price agreement with resale price p and discounted price $p - d(p)$. The discount $d(p) \in [0, p]$ is defined by*

$$d(p) := pF(p) - rF(\underline{p}) - \int_{\underline{p}}^{p} \psi_G^{-1} \circ \psi_F(x) f(x) dx.$$

The optimal delegation of resale limits the two inefficiencies that arise in a *laissez-faire* setting. First, it reduces the ability of the intermediary to exercise market power on the intensive margin. Indeed, intermediaries optimally choose resale prices below the monopoly price: $p(\theta_1) = \psi_F^{-1}(\psi_G(\theta_1)) \leq \psi_F^{-1}(\theta_1) = p^m(\theta_1)$. Second, it never excludes high-value consumers. Under *laissez-faire*, such consumers may not purchase the good when facing a low-value intermediary because the intermediary, wary of having to keep the good he does not value, does not purchase in the first place. In the optimal agreement, the seller circumvents this problem: when the intermediary has low value, he buys the good knowing he can sell it back at zero expected profits. Thus, the seller is able to implement the optimal auction through delegated contracting.

Optimal resale without buybacks. In some settings, it may be difficult for the seller to commit to a buyback policy. When the seller is a firm issuing securities, for example, a buyback policy may be seen as a collusive tactic between the seller and the intermediary—typically, a security underwriter—or a mechanism for the firm to inflate its own price. To capture these environments, we consider this natural restriction: what can the seller achieve when buyback policies are forbidden?

The role of buyback policies in the optimal mechanism is to allow the seller to leverage low-value intermediaries to reach high-value consumers. Without buybacks, this is impossible, leading to a tradeoff between extracting surplus and encouraging intermediation. Either low-value intermediaries must be incentivized to keep the good, limiting the seller's ability to exert monopoly power, or high-value consumers must be excluded, limiting the total surplus the seller can extract.

Our final result in this section shows it is never optimal to exclude high-value consumers, so all intermediaries buy the good. In general, the value of intermediation is higher than the monopoly profits the seller can extract by excluding some intermediaries. Further, we bound the seller's loss from being unable to implement buybacks. Let π^{BB} be the revenue the seller makes in the optimal delegation with buybacks, and π^{N} what he can make without buybacks.

Proposition 3. *Define $\hat{p} := \psi_F^{-1} \circ \psi_G(0)$. The set of contracts $\mathcal{C} = \{C^p\}_{p \in [\hat{p},1]}$ maximizes the revenue of the seller. Each contract $C^p \in \mathcal{C}$ is a price agreement with resale price p and discounted price $p - d(p)$ with discount $d(p) \in [0, p]$ defined by:*

$$d(p) := pF(p) - \int_{\hat{p}}^{p} \psi_G^{-1} \circ \psi_F(x) f(x) dx.$$

Moreover, $\frac{1}{2}\pi^{BB} \leq \pi^{N} \leq \pi^{BB}$.

3.3 Efficiency through Delegation

This section investigates when efficient outcomes can be implemented through delegated contracting in an allocation problem. As is standard, the designer wants to allocate a good to the party who values the good the most. By varying the initial distribution of property rights, this setting encompasses the bilateral trade problem in Myerson and Satterthwaite (1983), in which property rights are fully concentraded in the hands of one party, as well as partnership dissolution a *a la* Cramton et al. (1987). For any distribution of property rights, we first show that there is a unique budget-balanced mechanism which induces efficiency through delegation. We then prove a positive result: if property rights are all in the hands of the designer, efficiency dissolution is possible.

We complement this with a negative result. Cramton et al. (1987) show that a designer can efficiently dissolve a partnership as long as the initial ownership shares of the partners are not too unequal. This result, however, relies on the existence of a benevolent mediator who not only sets the rules of the interaction between the two partners, but also runs the bidding game that implements efficiency. We show below that the mediator is indispensable: the designer can never achieve efficiency if she cannot control the communication between the parties, regardless of the distributions of ownership and valuations.

We study a generalization of the standard resource allocation problem that accommodates both rival and non-rival resources (e.g., private and public goods), and assumes that the principal cannot subsidize the interaction between the players. As in centralized mechanism design, the prospects for achieving efficiency differ sharply across these cases. We first establish general conditions under which efficiency can be achieved when outside options are normalized to zero. We then characterize the highest outside options compatible with efficient implementation.

We focus on a quasilinear environment with symmetric players. Formally, for $i \in \{1,2\}$, types are $\theta_i \in [0,1]$, drawn independently according to F with full support and density f. The outcome space includes allocations $x \in X$, with X compact, and transfers $t_i \in \mathbb{R}$, $i \in \{1,2\}$, satisfying budget balance: $t_1 + t_2 = 0$.[5] Player i's utility under allocation x and transfer t_i is given by $u(x,\theta_i) - t_i$, where u is continuous in both arguments, and it is continuously differentiable and strictly increasing in θ_i for each $x \in X$.

An efficient allocation $x(\theta_1,\theta_2)$ exists for every type profile $(\theta_1,\theta_2) \in [0,1]^2$ and is characterized as

$$x(\theta_1,\theta_2) = \arg\max_{x \in X} \sum_{i=1}^{2} u(x,\theta_i),$$

which we assume to be uniquely defined.

To shorten notation, define $u_i[\theta_1,\theta_2] := u(x(\theta_1,\theta_2),\theta_i)$, and let the ex-post surplus be $s(\theta_1,\theta_2) := \sum_{i=1}^{2} u_i[\theta_1,\theta_2]$. The interim expected surplus conditional on θ_1 is

$$S(\theta_1) := \int_{\theta_2} s(\theta_1,\theta_2) F(d\theta_2),$$

and the ex-ante expected surplus is

$$\overline{S} := \int_{\theta_1} S(\theta_1) F(d\theta_1).$$

Outside options for players are given by $\underline{u}_i(\theta_i) := u(o,\theta_i)$, where o denotes the disagreement outcome. We impose the normalization that the lowest types have zero surplus and outside options:

$$\underline{u}_i(0) = 0 = s(0,0), \quad \text{for each } i \in \{1,2\}.$$

We distinguish between rival and non-rival resources. Resources are *(strictly) rival* if the surplus function s is (strictly) submodular, and *(strictly) non-rival* if s is (strictly) supermodular. For some intuition, consider a standard problem of providing a public good to the parties that costs c. Efficiency requires the principal to provide the good if and only if the sum of the players' types is above the cost, and the resulting surplus is $\max\{\theta_1 + \theta_2 - c, 0\}$, which is supermodular.

Our first result establishes that delegation can always achieve efficiency in rival allo-

[5] Our characterization will apply equally to weak budget balance, which requires $t_1 + t_2 \geq 0$

cation problems, but never in non-rival problems.

Proposition 4. *Suppose that $\underline{u}_i(\theta_i) = 0$ for all $\theta_i \in [0,1]$ and $i \in \{1,2\}$. Then:*

(i) If resources are (strictly) rival, the efficient allocation is implementable through delegation;

(ii) If resources are strictly non-rival, the efficient allocation is not implementable through delegation.

Proposition 4 mirrors well-known results from centralized mechanism design in the absence of outside options: private goods can be efficiently allocated under budget balance, whereas public goods cannot (Mailath and Postlewaite, 1990). This parallel underscores a positive message: Despite the additional frictions owing to loss of control, delegation does not fundamentally preclude allocative efficiency in rival environments. However, as we show later, the level of outside options is key to this result, and the class of environments where efficiency is achievable narrows under delegated contracting.

The proof of Proposition 4 proceeds by indirect construction. Leveraging Theorem 1, we show there are unique transfers that enable delegated implementation of the efficient allocation for rival goods. In contrast, for non-rival resources, we show that no direct mechanism can implement the efficient allocation while respecting the agent's DSIC and EPIR constraints.

Below, we describe contractual rights that support efficient outcomes through delegation, based on the construction used in the proof. Invoking budget balance, we can eliminate one of the transfers and describe an outcome by a pair (x, t_2): the allocation $x \in X$ and a transfer for the agent t_2. Consider the contractual rights $\mathcal{C} = \left[C^{\theta_1}\right]_{\theta_1 \in [0,1]}$, where each allowable menu is of the form $C^{\theta_1} = [x(\theta_1,\theta_2), t_2(\theta_1,\theta_2)]_{\theta_2 \in [0,1]}$ where:

$$t_2(\theta_1,\theta_2) = u_2[\theta_1,\theta_2] - s(\theta_1,\theta_2) + S(\theta_1) - S(0).$$

By budget balance, the delegate receives $-t_2(\theta_1,\theta_2)$. In equilibrium, we show it is optimal for the delegate of type θ_1 to choose the menu C^{θ_1}, and for the agent of type θ_2 to then choose the outcome associated with his type: $(x(\theta_1,\theta_2), t_2(\theta_1,\theta_2))$. These contractual rights therefore implement the efficient allocation.

Under truthful reporting by both players, the resulting ex-post payoffs for the agent and the delegate are, respectively:

$$U_a(\theta_1,\theta_2) = s(\theta_1,\theta_2) - S(\theta_1) + S(0),$$
$$U_d(\theta_1,\theta_2) = S(\theta_1) - S(0).$$

Several observations are worth highlighting. First, the structure of transfers reflects the dominant-strategy incentive compatibility constraint faced by the agent: transfers must replicate VCG payments up to a constant, which is determined by individual rationality. Second, and more surprisingly, the combination of the agent's DSIC constraint with the budget balance condition fully insures the delegate. That is, the delegate's ex-post payoff depends only on his own information and not on the agent's report. This outcome emerges because the agent, acting at the final stage, must fully internalize his effect on the surplus so that he is incentivized to maximize ex-post surplus. But that means that he must be the residual claimant in the interaction, while the delegate keeps the same amount of the surplus in all cases.

Finally, we explore the role of outside options, which were previously normalized to zero. Any delegated procedure that implements the efficient outcome must give the players the payoffs U_a and U_d above. Notably, budget balance and the agent's incentive constraints jointly leave no undetermined constants remaining. We can then use those payoffs to describe the maximum outside option of each player that guarantees efficiency, as described in the following proposition.

Proposition 5. *Suppose resources are rival. The efficient allocation is implementable through delegation if and only if:*

$$\underline{u}_1(\theta_1) \leq S(\theta_1) - S(0) \quad \forall \theta_1 \in [0,1], \tag{1}$$

$$\underline{u}_2(\theta_2) \leq \min_{\theta_1 \in [0,1]} \{s(\theta_1, \theta_2) - S(\theta_1) + S(0)\} \quad \forall \theta_2 \in [0,1]. \tag{2}$$

The bounds derived in Proposition 5 are asymmetric despite the two players being otherwise identical. This asymmetry arises because our mechanism leaves the agent exposed to the delegate's type but not vice versa. Consequently, the delegate's outside option is bounded by his insured interim payoff, whereas the agent's outside option is bounded by the minimum residual payoff across possible delegate types. This suggests that when outside options are relevant, optimally delegating authority entails both the design of contractual rights and whom to allocate those rights to. We pursue this direction in the next application, leveraging the characterization in Proposition 5.

3.4 Partnership Dissolution

Finally, we study the canonical partnership dissolution problem. In a seminal paper, Cramton et al. (1987) show that a designer can efficiently dissolve a partnership as long as the initial ownership shares of the partners are not too unequal. This result, however, relies on the existence of a benevolent mediator who not only sets the rules for the interaction between the two partners, but also runs the bidding game that implements efficiency. We show below that the mediator is indispensable: the designer can never achieve efficiency if she cannot control the communication between the parties, regardless of the distribution of ownership.

Our setup primarily follows that of Cramton et al. (1987). Two partners, $i \in \{1,2\}$ own $r_i \geq 0$ shares of an asset, with $r_1 + r_2 = 1$. Each partner's private valuation for the asset, $\theta_i \in [0,1]$, is independently drawn from the distribution F. We assume F is continuous and has full support with density f and mean $\mu > 0$. An allocation is a final share of the asset $q_i \geq 0$ kept by agent i, which generates a payoff to i of $u(q_i,\theta)$. We assume the payoff function u is continuous, convex in q, increasing in (q,θ_i), with $u(0,\theta_i) = 0$ and $u(1,\theta_i) = \theta_i$ for all $\theta_i \in [0,1]$. In contrast, in the standard partnership dissolution problem, $u(q,\theta_i) = \theta_i q$, so that players value each share of the asset equally. In practice, there may be gains from control and losses from diffuse ownership—that is, $u(q,\theta_i) \leq \theta_i q$—capturing the potential benefits of increased decision-making authority inside the firm. Similarly, the partners may discount the value of partial ownership. The outside option of player i is the value they obtain by keeping their initial ownership of the asset, that is $u(r_i,\theta_i)$.

The principal's goal is to design contractual rights that lead the asset to be allocated efficiently. Under the assumptions above, the efficient allocation gives the entire asset to the partner who values it most: $q_1(\theta_1,\theta_2) = \mathbb{1}_{\theta_1 \geq \theta_2}$. Because we will consider all possible combinations of initial ownership, (r_1,r_2), it is without loss of generality to assume the delegate is partner $i = 1$.

In this setting, the ex-post efficient surplus is $s(\theta_1,\theta_2) = \max\{\theta_1,\theta_2\}$ and, thus:

$$\begin{aligned} S(\theta_1) &= \int_{\theta_2} \max\{\theta_1,\theta_2\} F(d\theta_2) \\ &= \int_0^{\theta_1} \max\{\theta_1 - \theta_2, 0\} + \int_{\theta_2} \theta_2 F(d\theta_2) \\ &= I_F(\theta_1) + \mu, \end{aligned}$$

where $I_F(x) = \int_0^x F(x)dx$.

The proof of Proposition 4 shows there is a unique direct mechanism that implements efficiency while satisfying budget balance and DSIC for the agent. The fundamental question is whether this candidate mechanism also satisfies the players' individual rationality constraints. For intuition, consider the following implementation of such mechanism through delegation. Let the contract space be $\mathcal{C} = \{C^\lambda\}_{\lambda\in[0,1]}$. By choosing a menu C^λ, the delegate commits to a takeover "bid" amount, b^λ, and a takeover "ask" amount, a^λ. The agent, then, has the right to either forfeit all their shares and receive a payment b^λ or purchase all the delegate's shares at price a^λ. For this contract space to be consistent with the candidate mechanism, we need:

$$b^\lambda = \lambda - I_F(\lambda), \text{ and}$$

$$a^\lambda = I_F(\lambda).$$

If efficiency is achieved, it is attained in an equilibrium of this implementation in which the delegate truthfully reveals his type—by choosing contract $\lambda = \theta_1$—, and the agent takes over the firm if and only $\theta_2 \geq \theta_1$. Notice that regardless of the type of the agent, the payoff of the delegate is the same:

$$U_d(\theta_1,\theta_2) = I_F(\theta_1).$$

The agent, on the other hand, is the residual claimant of the ex-post surplus:

$$U_a(\theta_1,\theta_2) = \max\{\theta_1,\theta_2\} - I_F(\theta_1).$$

No gains from control: the impossibility of efficiency. We first consider the case in which there are no gains and losses from control, that is $u(q,\theta_i) = \theta_i q$ as in Cramton et al. (1987). In this setting, the delegate's interim participation constraint plays an important role. Bayesian incentive compatibility alone implies that his interim payoff must be the expected surplus up to a constant. But in the unique candidate mechanism, the constant is such that the worst type of the delegate receives no rents. The rationale for the zero payoff at the bottom is as follows. First, by efficiency, it must be that the delegate with type 0 always sells the asset, because he is never the efficient holder. Thus, all the types of the agent buy the asset when the delegate has type 0. But DSIC means that the agent essentially faces a posted price and, for all the agents to buy, the price must be zero. Therefore, DSIC and budget balance together imply that the type-0 delegate makes no profit: he always sells the good for free.

The fact that the delegate with the lowest type has zero rents is consequential. His payoff in the mechanism must grow by the probability of ultimately owning the asset, $Pr\{\theta_1 \geq \theta_2|\theta_1\} = F(\theta_1)$, following the envelope characterization of incentive compatibility. However, as his type grows, his outside option grows by r_1, and $r_1 > F(\theta_1)$

for sufficiently small θ_1. It follows that the delegate does not want to participate in the mechanism when his type is close to zero: he knows he will be expropriated for minimal payment, and prefers to stick with his initial ownership stake. The next result formalizes this discussion.

Proposition 6. *Assume there are no gains from control: $u(q,\theta_i) = \theta_i q$. Then, the partnership cannot be dissolved efficiently for any initial ownership distribution (r_1,r_2).*

Gains from control: whom to delegate to? Given that efficient dissolution is impossible when the partners value their initial stakes proportionally, we now study the case in which control can be beneficial to the partners: $u(q,\theta_i) \leq \theta_i q$. For an example, consider the case in which θ_i reflects the gains player i accrues from controlling the firm, but his ability to exercise any control requires at least 50% of the shares of the firm and grows with ownership. In that case, we could have $u(q,\theta_i) = \theta_i v(q)$, where $v(q) = 0$ for $q \leq .5$, meaning that the player enjoys no benefit for control if they are not the majority asset holder. If v is continuous, increasing, convex, and $v(1) = 1$, then this setting fits into our general partnership dissolution environment. Because ownership is initially dispersed, the status quo is less valuable for the partner who enjoys no gains from control, implying that the players' reservation utilities are lower. This reduction in outside options may make an efficient allocation possible. However, due to the asymmetries present in the efficient mechanism, the choice of delegate matters: efficiency may only be achievable when one of the two players holds contractual authority.

Our goal in this section is to ask who the delegate should be when the distribution of initial ownership is unequal. Without loss of generality, we assume $r_1 \geq r_2$. A naive intuition suggests that the partner with the highest outside option should be the delegate in the interaction. Indeed, we have seen that in the candidate efficient mechanism, the delegate is fully insured against all types of the agent, suggesting his outside option might be easier to satisfy; moreover, we know that a delegated mechanism must satisfy ex-post individual rationality for the agent, which indicates once more that it is easier to convince the delegate to join the mechanism than to persuade the agent to do so. This logic turns out to be incorrect: whenever the answer of whom to delegate to is uncontroversial—in the sense that it does not depend on the realization of players' types—it must be that the delegate is the player with the lowest outside option. The intuition follows the same logic as the impossibility of efficiency in Proposition 6: the most constrained player is the delegate with a low value, who knows he will likely be expropriated, but cannot receive any meaningful compensation for trading his shares.

Proposition 7. *Let* $\mu \geq \frac{1}{2}$ *and* $r_1 \geq r_2$. *If efficiency can be achieved by delegating to player 1, then it can also be achieved by delegating to player 2.*

References

Alonso, R., and Matouschek, N. (2008). Optimal Delegation. *The Review of Economic Studies*, *75*(1), 259–293.

Amador, M., and Bagwell, K. (2013). The Theory of Optimal Delegation With an Application to Tariff Caps. *Econometrica*, *81*(4), 1541–1599.

Amador, M., and Bagwell, K. (2022). Regulating a monopolist with uncertain costs without transfers. *Theoretical Economics*, *17*(4), 1719–1760.

Ambrus, A., and Egorov, G. (2017). Delegation and nonmonetary incentives. *Journal of Economic Theory*, *171*, 101–135.

Baliga, S., and Sjöström, T. (1998). Decentralization and Collusion. *Journal of Economic Theory*, *83*(2), 196–232.

Bhaskar, D., McClellan, A., and Sadler, E. (2023). Regulation Design in Insurance Markets. *American Economic Review*, *113*(10), 2546–2580.

Cramton, P., Gibbons, R., and Klemperer, P. (1987). Dissolving a Partnership Efficiently. *Econometrica*, *55*(3), 615.

Dworczak, P., and Muir, E. (2024). *A Mechanism-Design Approach to Property Rights* (Unpublished Manuscript).

Faure-Grimaud, A., and Martimort, D. (2001). On some agency costs of intermediated contracting. *Economics Letters*, *71*(1), 75–82.

Frankel, A. (2014). Aligned Delegation. *American Economic Review*, *104*(1), 66–83.

Guo, Y. (2016). Dynamic Delegation of Experimentation. *American Economic Review*, *106*(8), 1969–2008.

Hiriart, Y., and Martimort, D. (2012). How much discretion for risk regulators? *The RAND Journal of Economics*, *43*(2), 283–314.

Holmström, B. (1979). Groves' scheme on restricted domains. *Econometrica: Journal of the Econometric Society*, 1137–1144.

Holmström, B. (1980). *On the theory of delegation* (Discussion Paper).

Kleiner, A., Moldovanu, B., and Strack, P. (2021). Extreme Points and Majorization: Economic Applications. *Econometrica*, *89*(4), 1557–1593.

Koessler, F., and Martimort, D. (2012). Optimal delegation with multi-dimensional decisions. *Journal of Economic Theory*, *147*(5), 1850–1881.

Kolotilin, A., and Zapechelnyuk, A. (2025). Persuasion Meets Delegation. *Econometrica*, *93*(1), 195–228.

Krishna, V., and Perry, M. (1998). Efficient mechanism design.

Kundu, T., and Nilssen, T. (2020). Delegation of Regulation. *The Journal of Industrial Economics*, *68*(3), 445–482.

Laffont, J.-J., and Rochet, J.-C. (1997). Collusion in Organizations. *The Scandinavian Journal of Economics*, *99*(4), 485–495.

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

Malladi, S. (2022). *Delegated Costly Screening* (Unpublished Manuscript).

Martimort, D., Pouyet, J., and Staropoli, C. (2020). Use and abuse of regulated prices in electricity markets: "How to regulate regulated prices?". *Journal of Economics & Management Strategy*, *29*(3), 605–634.

Maskin, E., and Tirole, J. (1990). The principal-agent relationship with an informed principal: The case of private values. *Econometrica: Journal of the Econometric Society*, 379–409.

McAfee, R. P., and McMillan, J. (1995). Organizational Diseconomies of Scale. *Journal of Economics & Management Strategy*, *4*(3), 399–426.

Melumad, N., Mookherjee, D., and Reichelstein, S. (1992). A theory of responsibility centers. *Journal of Accounting and Economics*, *15*(4), 445–484.

Melumad, N. D., Mookherjee, D., and Reichelstein, S. (1995). Hierarchical Decentralization of Incentive Contracts. *The RAND Journal of Economics*, *26*(4), 654.

Melumad, N. D., and Shibano, T. (1991). Communication in Settings with No Transfers. *The RAND Journal of Economics*, *22*(2), 173.

Mitchell, M. (2025). *Pricing with Repeated Resale: Can Creators Benefit From Perpetual Royalties* (Unpublished Manuscript).

Mookherjee, D. (2006). Decentralization, Hierarchies, and Incentives: A Mechanism Design Perspective. *Journal of Economic Literature*, *44*(2), 367–390.

Mookherjee, D. (2012). Incentives in Hierarchies. In *The Handbook of Organizational Economics* (pp. 764–798). Princeton University Press.

Mookherjee, D., and Tsumagari, M. (2004). The Organization of Supplier Networks: Effects of Delegation and Intermediation. *Econometrica*, *72*(4), 1179–1219.

Myerson, R. B. (1981). Optimal Auction Design. *Mathematics of Operations Research*, *6*(1), 58–73.

Myerson, R. B., and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, *29*(2), 265–281.

Severinov, S. (2003). *Optimal organization: Centralization, decentralization or delegation?* (Unpublished Manuscript).

Tirole, J. (1986). Hierarchies and Bureaucracies: On the Role of Collusion in Organizations. *Journal of Law, Economics, & Organization*, *2*(2), 181–214.

Tirole, J. (1992). Collusion and the Theory of Organizations. *Advances in economic theory*, *2*, 151–206.

Zheng, C. Z. (2002). Optimal Auction with Resale. *Econometrica*, *70*(6), 2197–2224.

Appendix: Proofs

Proof of Theorem 1

Only if ⇒

Assume that Y is implementable through delegation. Let $\mathcal{C}$ be the contract space, and $(\sigma_1, \{\sigma_2^C, \mu^C\}_{C\in\mathcal{C}})$ an associated equilibrium that implements Y. Because any Perfect Bayesian Equilibrium is a Bayes-Nash Equilibrium, the revelation principle implies that there exists a direct centralized mechanism in which Y is implemented by truthful revelation. Formally, there exists $\mathcal{M} = (M, g)$ such that $M = \Theta_1 \times \Theta_2$, with optimal strategies $s_i(\theta_i) = \theta_i$, for $i \in \{1,2\}$. Because this equilibrium implements Y, it must be the case that $g(\theta_1, \theta_2) = Y(\theta_1, \theta_2)$ for all $(\theta_1, \theta_2) \in \operatorname{supp} \mu_o$. Moreover, because Y is deterministic, $\sigma_2^C(\theta_2)$ is deterministic for all $C \in \operatorname{supp} \sigma_1(\theta_1)$ and because Y is implemented through delegation:

$$g(\theta_1, \theta_2) = \sigma_2^C(\theta_2), \tag{3}$$

for all $(\theta_1, \theta_2) \in \operatorname{supp} \mu_o$, $C \in \operatorname{supp} \sigma_1(\theta_1)$.

Because Y is implemented by $\mathcal{M}$, it must satisfy BIC and IIR for both players. In what follows we prove it also satisfies (i) DSIC and (ii) EPIR for the agent.

Fix any $\theta_1 \in \Theta_1$, and $C \in \operatorname{supp} \sigma_1(\theta_1)$. Because $(\sigma_1, \{\sigma_2^C, \mu^C\}_{C\in\mathcal{C}})$ was an equilibrium in the delegated contracting game, we have, for any agent's information θ_2 such that $(\theta_1, \theta_2) \in \operatorname{supp} \mu_o$ and alternative outcome available to the agent, $x' \in C$:

$$u_2\left(\sigma_2^C(\theta_2), \theta_2\right) \geq u_2\left(x', \theta_2\right).$$

Note that μ^C does not enter this equilibrium constraint since θ_1 is payoff-irrelevant for the agent. We can, in particular, choose $x' = \sigma_2^C(\theta_2')$ for any type of the agent with $(\theta_1, \theta_2') \in \operatorname{supp} \mu_o$, and take expectation over C on both sides with respect to $\sigma_1(\theta_1)$ to obtain:

$$\begin{aligned} u_2\left(g(\theta_1, \theta_2), \theta_2\right) &= \int_C u_2\left(\sigma_2^C(\theta_2), \theta_2\right) \sigma_1(\theta_1)(dC) \\ &\geq \int_C u_2\left(\sigma_2^C(\theta_2'), \theta_2\right) \sigma_1(\theta_1)(dC) = u_2(g(\theta_1, \theta_2'), \theta_2), \end{aligned}$$

where the first and last equalities come from the equality in 3. Because θ_1 was arbitrary, the inequality above shows truthful revelation is a dominant strategy for the agent in direct mechanism $\mathcal{M}$.

Because we assume the agent can always choose o, the result above also implies, for any θ_1:

$$u_2(g(\theta_1,\theta_2),\theta_2) \geq u_2(o,\theta_2),$$

so that truthful revelation also satisfies ex-post individual rationality. Because, by implementation, $g = Y$, we obtain that Y satisfies (i) and (ii).

If $\Leftarrow$

To prove the converse, let Y be implementable in a centralized mechanism, and let (i) and (ii) hold. By the revelation principle, there is a direct mechanism $\mathcal{M} = (M,g)$ that implements Y in a truthful-revealing equilibrium. Because it implements Y we must have:

$$g(\theta_1,\theta_2) = Y(\theta_1,\theta_2). \tag{4}$$

Consider the contract space $\mathcal{C} = [C^{\theta_1}]_{\theta_1 \in \Theta_1}$, where

$$C^{\theta_1} = \{y : \exists \theta_2 \text{ with } (\theta_1,\theta_2) \in \operatorname{supp} \mu_o \text{ and } y = Y(\theta_1,\theta_2)\}.$$

Given this contract space, We will show that it is a PBE for the delegation game for each type of the delegate to offer menu C^{θ_1}, and, upon observing C_1^{θ}, for the agent to choose contract $Y(\theta_1,\theta_2) \in C^{\theta_1}$.

Starting with the agent's best response, because of (i), and because $\mathcal{M}$ implements Y, for all $\theta_1 \in \Theta_1$, $\theta_2,\theta_2' \in \Theta_2$:

$$\begin{aligned} u_2(Y(\theta_1,\theta_2),\theta_2) &= u_2(g(\theta_1,\theta_2),\theta_2) \\ \geq u_2(g(\theta_1,\theta_2'),\theta_2) &= u_2(Y(\theta_1,\theta_2'),\theta_2), \end{aligned}$$

where the equalities follow from equation 4.

Because of (ii), and because $\mathcal{M}$ implements Y for all $\theta_1 \in \Theta_1$, $\theta_2,\theta_2' \in \Theta_2$:

$$u_2(Y(\theta_1,\theta_2),\theta_2) = u_2(g(\theta_1,\theta_2),\theta_2) \geq u_2(o,\theta_2),$$

so the two inequalities above show that it is optimal for the agent to reveal his type truthfully in the delegation game for any contract C^{θ_1}, for any beliefs the agent might hold about the delegate.

Next, because $\mathcal{M}$ implements Y in BIC for the delegate:

$$\int_{\theta_2} u_1(Y(\theta_1,\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2) = \int_{\theta_2} u_1(g(\theta_1,\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2)$$
$$\geq \int_{\theta_2} u_1(g(\theta_1',\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2) = u_1(Y(\theta_1',\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2),$$

Moreover, because of IIR:

$$\int_{\theta_2} u_1(Y(\theta_1,\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2) = \int_{\theta_2} u_1(g(\theta_1,\theta_2),\theta_1)\mu_o(\theta_1,d\theta_2)$$
$$\geq u_1(o,\theta_1),$$

implying it is an equilibrium of the delegated game for the delegate to choose the contract associated to his type, regardless of the agent's beliefs.

Given that it is optimal for both players to truthfully reveal their types in the delegated game, the outcome implemented in this equilibrium is $Y(\theta_1,\theta_2)$, finishing the proof. ■

Proof of Proposition 1

We now provide a more formal proof leveraging Theorem 1 and a relaxation of the principal's problem.

The principal's full problem is the following.

$$\begin{aligned} &\max_{q,t} \mathbb{E}_{b,s}[bq(b,s) - (1+\alpha)t(b,s)] \\ s.t. \quad & \text{Buyer BIC} \\ & \text{Buyer IIR} \\ & \text{Seller DSIC} \\ & \text{Seller EPIR} \end{aligned}$$

We move towards a relaxation using the type-b delegate's unconstrained screening problem.

$$\begin{aligned} \max_{q(b,\cdot),t(b,\cdot)} \quad & \mathbb{E}_s[bq(b,s) - t(b,s)] \\ s.t. \quad & \text{Agent IC} \\ & \text{Agent IR} \end{aligned}$$

Let $U^*(b)$ denote the value of the above problem for each b, and let $Q^*(b) = \mathbb{E}_s[q^*(b,s)]$ denote the expected output under the b-optimal contract. Note that $Q^*(b)$ is weakly increasing in b, and $U^*(b)$ is weakly increasing and convex, with $U^*(0) = 0$.

Since the principal can only restrict the contract space, DSIC and EPIR of the agent imply that the delegate can never achieve a payoff above $U^*(b)$. Hence, we consider the following relaxed problem, where the principal solves for expected quantity $Q(b)$ and transfer $T(b)$ associated with the type-b delegate:

$$\begin{aligned} \max_{Q(b),T(b)} \quad & \mathbb{E}_b[bQ(b) - T(b)] \\ s.t. \quad & \text{Delegate BIC} \\ & \text{Delegate IIR} \\ & U(b) \le U^*(b) \quad \forall b \end{aligned}$$

We now proceed with the usual Myersonian transformation. Standard substitution of the envelope and monotonicity conditions yield the following program.

$$\begin{aligned} \max_{Q(b)} \quad & \int_0^1 \phi(b) f(b) Q(b)\, db \\ \text{s.t.} \quad & Q(b) \ge 0 \quad \text{for all } b \in [0,1] \\ & Q(b) \text{ non-decreasing on } [0,1] \\ & \int_0^b Q(r)\, dr \le \int_0^b Q^*(r)\, dr \quad \text{for all } b \in [0,1] \end{aligned}$$

Here, $\phi(b) := -\alpha b + (1+\alpha)\frac{1-F(b)}{f(b)}$, and f is a strictly positive density and differentiable on $[0,1]$ with associated CDF F. We assume the monotone hazard rate (MHR) condition: $(1 - F(b))/f(b)$ is decreasing in b. The function $Q^*(b)$ is a non-negative, non-decreasing benchmark allocation rule.

We define the integrand $J(b) := \phi(b)f(b) = -\alpha b f(b) + (1+\alpha)(1-F(b))$. We now establish the structure of the optimal solution in four steps.

Step 1: $J(b)$ crosses zero exactly once

We first show that $J(b)$ is positive for small b and negative for large b:

$$J(0) = (1+\alpha)(1-F(0)) = 1+\alpha > 0$$
$$J(1) = -\alpha f(1) + (1+\alpha)(1-F(1)) = -\alpha f(1) < 0$$

By continuity of J, there exists $\bar{b} \in (0,1)$ such that $J(\bar{b}) = 0$. Furthermore, under the MHR assumption, $\phi(b)$ is strictly decreasing. Since $f(b) > 0$, $J(b)$ has the same sign as $\phi(b)$ and crosses zero exactly once from above at $\bar{b}$:

$$J(b) > 0 \text{ for } b < \bar{b}, \quad J(b) < 0 \text{ for } b > \bar{b}$$

Step 2: Optimally, $Q(b)$ is constant on $[\bar{b},1]$

Take any feasible Q for which there exist $b_1, b_2 \in [\bar{b},1)$, with $Q(b_1) < Q(b_2)$. By monotonicity of Q, $b_2 > b_1$ and $Q(b_1) < Q(b)$ for all $b \geq b_2$. Define a new function $\tilde{Q}$ that is identical to Q on $[0,b_1]$ and constant thereafter:

$$\tilde{Q}(b) = \begin{cases} Q(b) & \text{if } b \leq b_1 \\ Q(b_1) & \text{if } b > b_1 \end{cases}$$

Then $\tilde{Q}$ is non-decreasing, satisfies the same cumulative constraint (as $\tilde{Q}(b) \leq Q(b)$ for all b), and yields a strictly higher objective value because $J(b) < 0$ on $[b_1,1]$ and $\tilde{Q}(b) < Q(b)$ there. Therefore Q cannot be optimal. We then conclude that any optimal Q must be constant on $[\bar{b},1]$.

Step 3: $J(b)$ is decreasing on $\{b : J(b) > 0\}$

We verify this directly using the expression $J(b) = -\alpha b f(b) + (1+\alpha)(1-F(b)) = \phi(b)f(b)$. Differentiating:

$$J'(b) = -f(b)(2\alpha+1) - \alpha b f'(b) = \phi(b)f'(b) + \phi'(b)f(b)$$

We show that the equation above implies $J'(b) < 0$ in two cases:

- If $f'(b) > 0$, then both terms in the first expression are negative
- If $f'(b) \leq 0$, then both terms in the second expression are negative because $\phi(b) > 0$ and $\phi'(b) < 0$ by the MHR property

Hence, $J(b)$ is strictly decreasing wherever it is positive.

Step 4: The optimal Q matches Q^* up to a cutoff $\hat{b} \leq \bar{b}$

Since $J(b)$ is strictly decreasing on the relevant region $[0, \bar{b}]$, the marginal value of increasing $Q(b)$ is highest at the lowest types. Given that Q^* is non-decreasing and the cumulative constraint is:

$$\int_0^b Q(r)\,dr \leq \int_0^b Q^*(r)\,dr,$$

the optimal solution must match Q^* starting from $b = 0$ and continue doing so until the constraint binds. This follows from a standard monotone reallocation argument as in Kleiner et al. (2021): when the integrand is strictly decreasing, the optimum is to match the upper bound allocation as fully as possible starting from the left.

Formally, the principal's problem can be decomposed into a two-step maximization problem: first deciding the maximal point at which Q is no longer increasing, then deciding the level of Q below this cutoff:

$$\begin{aligned}
\max_{Q(b),\hat{b}} \quad & \int_0^{\hat{b}} J(b)Q(b)\,db + \int_{\hat{b}}^1 J(b)Q(\hat{b})\,db \\
\text{s.t.} \quad & Q(b) \geq 0 \quad \text{for all } b \in [0,1] \\
& Q(b) \text{ non-decreasing on } [0,\hat{b}] \\
& \int_0^b Q(r)\,dr \leq \int_0^b Q^*(r)\,dr \quad \text{for all } b \in [0,\hat{b}] \\
& \int_0^{\hat{b}} Q(r)\,dr = \int_0^{\hat{b}} Q^*(r)\,dr
\end{aligned}$$

For any given $\hat{b}$, the optimal $Q(b)$ follows from Proposition 3 in Kleiner et al. (2021). Therefore, there exists a cutoff $\hat{b} \leq \bar{b}$ such that:

$$Q(b) = Q^*(b) \quad \text{for all } b \leq \hat{b},$$

and the cumulative constraint binds at $\hat{b}$:

$$\int_0^{\hat{b}} Q(r)\,dr = \int_0^{\hat{b}} Q^*(r)\,dr$$

Thereafter, $Q(b)$ remains constant at $Q^*(\hat{b})$ for all $b > \hat{b}$.

Clearly, this payoff can be implemented in the original problem by allowing each type of delegate to choose their optimal screening contract up to the level $\hat{b}$, at which point the quantity no longer changes. ■

Proof of Proposition 2

By Theorem 1, it is clear that if a contract space allows the seller to achieve the same revenue as the optimal BIC auction, it must be optimal. Indeed, we know that the optimal auction can be implemented in dominant strategies, so Theorem 1 tells us there exists a contract space that attains the same revenue. We formalize the indirect contracts in the statement of the proposition that achieve this goal.

Define $\underline{p} = \psi_F^{-1}(0)$, $r := \psi_G^{-1}(0)$, and $d(p)$ as in the statement. The contract space is $\mathcal{C} = \{C^p\}_{p\in[\underline{p},1]}$ such that the message space for each contract C^p is $M = \{0,1\}$, where 0 should be understood as the consumer not buying and 1 as the consumer buying at price p. Moreover, the outcomes of any contract p are given by (q_i^p, t_i^p), where q_i^p is the probability that agent i keeps the good and t_i^p is the net outflow of money from that agent. The outcome functions satisfy:

$$q_2^p(m) = \mathbb{1}_{m=1} \qquad\qquad q_1^p(m) = \mathbb{1}_{m=0,p>\underline{p}}$$

$$t_2^p(m) = p\mathbb{1}_{m=1} \qquad\qquad t_1^p(m) = (p - d(p)) - p\mathbb{1}_{m=1} - r\mathbb{1}_{m=0,p=\underline{p}}.$$

In words, the buyer purchases the good whenever his message is $m = 1$, in which case he pays the resale price of the appropriate contract, p. The intermediary keeps the good whenever the message is $m = 0$, except if $p = \underline{p}$, in which case he returns the product if $m = 0$. The intermediary's transfer has three possible cases. The outflow is always the purchase price $p - d(p)$ at which the intermediary buys from the seller. If $m = 1$, the intermediary sells at p, resulting in a net inflow of $d(p)$; if $m = 0$, and the contract is $\underline{p}$, the intermediary returns the good in case $m = 0$ and gets the refund for a net inflow of r.

We now solve for the unique (up to measure zero) equilibrium induced by this game.

It is clear that it is optimal for the final consumer to buy ($m = 1$) if and only if $p \geq \theta_2$. For the optimal choice of price by the intermediary, notice that he maximizes:

$$\max_p p(1 - F(p)) + \theta_1 F(p) - (p - d(p)) = \max_p (\theta_1 - p)F(p) + d(p).$$

This objective is differentiable, and its derivative is:

$$(\theta_1 - p)f(p) - F(p) + F(p) + pf(p) - \psi_G^{-1} \circ \psi_F(p)f(p) = \left(\theta_1 - \psi_G^{-1} \circ \psi_F(p)\right) f(p),$$

which is positive for $p \leq \psi_F^{-1} \circ \psi_G(\theta_1)$, and negative otherwise. Therefore, the optimal price chosen by an intermediary of type θ_1 is $p(\theta_1) := \psi_F^{-1} \circ \psi_G(\theta_1)$.

Because the minimal allowed price is $\underline{p} = \psi_F^{-1}(0)$, it must be that all types $\theta_1 \leq \psi_G^{-1}(0)$ choose $\underline{p}$. Moreover, for $\theta_1 := \psi_G^{-1}(0)$, two observations are important. (1) This intermediary is indifferent between returning and not returning the good, implying all types of the intermediary $\theta_1 < \psi_G^{-1}(0)$ return the product when they cannot sell it. (2) computing the payoff of type $\theta_1 = \psi_G^{-1}(0) = r$:

$$(r - \underline{p})F(\underline{p}) + d(\underline{p}) = 0,$$

so $\psi_G^{-1}(0)$ is held to his outside option. Thus, the scheme we designed generates the same allocations as the optimal BIC auction, and gives the bidders their same interim payoffs. As a conclusion, it has the same revenue as the optimal BIC auction and is, therefore, optimal.

Finally, we show that $d(p) \in [0, p]$. That $d(p) \leq p$ is evident by its formula. To show that $d(p) \geq 0$, we start by noticing that $\psi_G \geq \psi_F$. Indeed, our gains-from-trade assumption is equivalent to:

$$\frac{G(x) - G(t)}{1 - G(t)} \geq \frac{F(x) - F(t)}{1 - F(t)},$$

for all $0 \leq t \leq x \leq 1$. Therefore:

$$\begin{aligned} 0 &\leq (G(x) - G(t))\,(1 - F(t)) - (F(x) - F(t))\,(1 - G(t)) \\ &= (1 - F(x))(1 - G(t)) - (1 - F(t))(1 - G(x)), \end{aligned}$$

and, therefore:

$$\frac{1 - F(x)}{1 - G(x)} \geq \frac{1 - F(t)}{1 - G(t)},$$

for all $0 \le t \le x \le 1$. In other words, the function $R(x) = \frac{1-F(x)}{1-G(x)}$ is increasing. Because it is also differentiable, we must have:

$$0 \le R'(x) = \frac{g(x)f(x)}{(1-F(x))^2}\left(\frac{1-F(x)}{f(x)} - \frac{1-G(x)}{g(x)}\right),$$

which concludes that $\psi_F(x) = x - \frac{1-F(x)}{f(x)} \le x - \frac{1-G(x)}{g(x)} = \psi_G(x)$ for all $x \in [0,1]$.

We now will use this fact to prove $d(p) \ge 0$. Note that d is differentiable with:

$$\begin{aligned} d'(p) &= \left(p - \psi_G^{-1} \circ \psi_F(p)\right) f(p) + F(p) \\ &= (p - \psi_F(p)) f(p) + F(p) + \left(\psi_F(p) - \psi_G^{-1} \circ \psi_F(p)\right) f(p) \\ &= 1 + \left(\psi_F(p) - \psi_G^{-1} \circ \psi_F(p)\right) f(p) \\ &\ge 1 + \left(\psi_F(p) - \psi_F^{-1} \circ \psi_F(p)\right) f(p) \\ &= 1 - (1 - F(p)) \ge 0, \end{aligned} \tag{5}$$

where the second equality adds and subtracts $\psi_F(p)f(p)$, the third equality uses the definition of ψ_F, and the inequality follows from $\psi_F \le \psi_G$, which implies $\psi_F^{-1} \ge \psi_G^{-1}$.

Because d is increasing, it is sufficient to check $d(\underline{p}) \ge 0$. For that:

$$d(\underline{p}) = (\underline{p} - r)F(\underline{p}) = \left(\psi_F^{-1}(0) - \psi_G^{-1}(0)\right) F(\underline{p}) \ge 0,$$

with the inequality following again from $\psi_F \le \psi_G$. ■

Proof of Proposition 3

Once more, we leverage Theorem 1. A direct standard mechanism is a tuple $(q_1, q_2, t_1, t_2) : [0,1]^2 \to \mathbb{R}^4$, with $q_i \in [0,1]$, being the probability that party i is allocated the good, with $q_1(\theta_1, \theta_2) + q_2(\theta_1, \theta_2) \le 1$; and t_i the outflow of cash from player i. The restriction of no buybacks can be rewritten as: $q_1(\theta_1, \theta_2) + q_2(\theta_1, \theta_2) = \gamma(\theta_1)$, for some function γ. That is, the total allocation of the good between the final consumer and the intermediary is determined solely by the type of the intermediary—and, therefore, at the time the intermediary contracts with the seller. By Theorem 1, The seller chooses q_1, q_2, t_1, t_2 and γ to maximize revenue subject to BIC, IIR for the intermediary and DSIC, EPIR for the final consumer.

We first argue that $\gamma(\theta_1) = 1$ is without loss of optimality. First, the seller does not value the product, so the only reason she would retain the good is for incentives. We shall show that, for any mechanism, there is an alternative mechanism with the same seller revenue such that there is no retention. Fix $(q_1, q_2, t_1, t_2, \gamma)$ satisfying all the constraints. Consider the following alternative $(q_1, q_2 + (1-\gamma), t_1, t_2, 1)$: that is, the mechanism is the same except that any units previously retained are now allocated to the buyer. The intermediary's allocation and transfers are exactly the same, so his incentives are maintained. The buyer's ex-post incentives are also unchanged: we are only adding to his allocation a constant given the type of the intermediary, which they observe. Moreover, individual rationality can only be slackened since now the mechanism is more generous to the buyer. Thus, buyers' incentives to participate and reveal his type are maintained as well. Therefore, the problem of the principal is to maximize revenue subject to BIC, IIR for the intermediary and DSIC, EPIR for the final consumer subject to no retention. Thus, the optimal mechanism is a virtual auction with no reserve price.

We now show how to implement such mechanism through delegation with price agreement contracts, following closely the proof of Proposition 2. The contract space is $\mathcal{C} = \{C^p\}_{p \in [0,1]}$ such that the message space for each contract C^p is $M = \{0, 1\}$, where 0 should be understood as the consumer not buying and 1 as the consumer buying. Moreover, the outcomes of any contract p are given by (q_i^p, t_i^p), where q_i^p is the probability that agent i keeps the good and t_i^p is the net outflow of money from that agent. The outcome functions satisfy:

$$\begin{aligned} q_2^p(m) &= \mathbb{1}_{m=1} & \qquad q_1^p(m) &= \mathbb{1}_{m=0} \\ t_2^p(m) &= p\mathbb{1}_{m=1} & \qquad t_1^p(m) &= (p - d(p)) - p\mathbb{1}_{m=1}. \end{aligned}$$

In words, the buyer keeps the good whenever his message is $m = 1$, in which case he pays the resale price of the appropriate contract, p. The intermediary keeps the good whenever the message is $m = 0$. The intermediary's transfer the purchase price $p - d(p)$ at which the intermediary buys from the seller, but if $m = 1$, the intermediary sells at p, resulting in a net inflow of $d(p)$.

We now solve for the unique (up to measure zero) equilibrium induced by this game. It is clear that it is optimal for the final consumer to buy ($m = 1$) if and only if $p \geq \theta_2$. For the optimal choice of price by the intermediary, notice that he maximizes:

$$\max_p p(1 - F(p)) + \theta_1 F(p) - (p - d(p)) = \max_p (\theta_1 - p)F(p) + d(p).$$

This objective is differentiable, and its derivative is:

$$(\theta_1 - p)f(p) - F(p) + F(p) + pf(p) - \psi_G^{-1} \circ \psi_F(p)f(p) = \left(\theta_1 - \psi_G^{-1} \circ \psi_F(p)\right) f(p),$$

which is positive for $p \leq \psi_F^{-1} \circ \psi_G(\theta_1)$, and negative otherwise. Therefore, the optimal price chosen by an intermediary of type θ_1 is $p(\theta_1) := \psi_F^{-1} \circ \psi_G(\theta_1)$. Moreover, an intermediary with $\theta_1 = 0$ chooses $p(0) = \hat{p}$ and makes profits:

$$(0 - \hat{p})\, F(\hat{p}) + \hat{p}F(\hat{p}) = 0.$$

Because this equilibrium has the same allocations as the virtual auction with no reserve, and the outside options of the lowest type players are the same, the seller must obtain the same revenue. Thus, it implements the optimal revenue.

Bounds on revenue

By the discussion in this proof and in the proof of Proposition 2, we obtain that the revenue with buybacks is $\pi^{BB} = \mathbb{E}\left[\max\{0, \psi_G(\theta_1), \psi_F(\theta_2)\}\right]$ and the revenue without buyback is $\pi^N = \mathbb{E}\left[\max\{\psi_G(\theta_1), \psi_F(\theta_2)\}\right]$. Thus:

$$\begin{aligned}
\pi^{BB} = \mathbb{E}\left[\max\{0, \psi_G(\theta_1), \psi_F(\theta_2)\}\right] &\leq \mathbb{E}\left[\max\{0, \psi_G(\theta_1)\} + \max\{0, \psi_F(\theta_2)\}\right] \\
&= \mathbb{E}\left[\max\{\mathbb{E}_{\theta_2}[\psi_F(\theta_2)], \psi_G(\theta_1)\} + \max\{\mathbb{E}_{\theta_1}[\psi_G(\theta_1)], \psi_F(\theta_2)\}\right] \\
&\leq \mathbb{E}\left[\max\{\psi_F(\theta_2), \psi_G(\theta_1)\} + \max\{\psi_G(\theta_1), \psi_F(\theta_2)\}\right] = 2\pi^N,
\end{aligned}$$

with the second equality following from the property that $\min \operatorname{supp} F = \min \operatorname{supp} G = 0$, which implies that the expected virtual value is zero; and the second inequality being a consequence of convexity of the maximum. ■

Proof of Proposition 4 and Proposition 5

We construct the unique direct standard mechanism that can implement efficiency while satisfying BIC for the delegate and DSIC for the agent. We then proceed to derive the bounds on outside options obtained in Proposition 5. Finally, we conclude by showing Proposition 4. By Theorem 1, efficiency can be implemented through delegation if and only if it can be implemented in a centralized mechanism and it is dominant-strategy incentive compatible and ex-post individually rational for the agent. Following Holmström (1979), dominant-strategy incentive compatibility implies the transfer paid by the agent must be the VCG transfer up to a constant:

$$t_2(\hat{\theta}_1, \hat{\theta}_2) = -u_1[\hat{\theta}_1, \hat{\theta}_2] + h(\hat{\theta}_1),$$

for some function h that does not depend on $\hat{\theta}_2$. By budget balance, $t_1 = -t_2$ so $t_1(\hat{\theta}_1, \hat{\theta}_2) = u_1[\hat{\theta}_1, \hat{\theta}_2] - h(\hat{\theta}_1)$. In expectation, the payments are $t_1(\hat{\theta}_1) := \mathbb{E}_{\tilde{\theta}_2}[t_1(\hat{\theta}_1, \tilde{\theta}_2)] = \mathbb{E}_{\tilde{\theta}_2}[u_1[\hat{\theta}_1, \tilde{\theta}_2] - h(\hat{\theta}_1)]$. Notice that the second term, $h(\hat{\theta}_1)$ is independent of $\hat{\theta}_2$, so it can be pulled outside of the expectation.

Again, since the BIC expected payments are unique up to a constant, BIC implies that the delegate must be compensated by the expected surplus of the agent up to a constant that does not depend on his report $t_1(\hat{\theta}_1) = \mathbb{E}_{\tilde{\theta}_2}[-u_2[\hat{\theta}_1, \tilde{\theta}_2]] + c$.

Combining the previous two equations, we have:

$$\mathbb{E}_{\tilde{\theta}_2}[u_1[\hat{\theta}_1, \tilde{\theta}_2]] - h(\hat{\theta}_1) = \mathbb{E}_{\tilde{\theta}_2}[-u_2[\hat{\theta}_1, \tilde{\theta}_2]] + c$$

So,

$$\begin{aligned} h(\hat{\theta}_1) &= \mathbb{E}_{\tilde{\theta}_2}[u_1[\hat{\theta}_1, \tilde{\theta}_2] + u_2[\hat{\theta}_1, \tilde{\theta}_2]] - c \\ &= S(\theta_1) - c \end{aligned}$$

Therefore, the payments are pinned down up to one constant, c.

$$\begin{aligned} t_2(\hat{\theta}_1, \hat{\theta}_2) &= -u_1[\hat{\theta}_1, \hat{\theta}_2] + S(\hat{\theta}_1) - c \\ t_1(\hat{\theta}_1, \hat{\theta}_2) &= -t_2(\hat{\theta}_1, \hat{\theta}_2) \\ &= u_1[\hat{\theta}_1, \hat{\theta}_2] - S(\hat{\theta}_1) + c \end{aligned}$$

Given these transfers, we can calculate the ex-post payoffs of the two players:

$$\begin{aligned} u_2(\theta_1, \theta_2) &= u_2[\theta_1, \theta_2] + u_1[\theta_1, \theta_2] - S(\theta_1) + c \\ &= s(\theta_1, \theta_2) - S(\theta_1) + c \\ u_1(\theta_1, \theta_2) &= S(\theta_1) - c \end{aligned}$$

Using the payoffs, we can now write the IR constraints

$$s(\theta_1, \theta_2) - S(\theta_1) + c \geq \underline{u}_2(\theta_2) \quad \forall \theta_1, \theta_2 \qquad \text{(Ex-post IR for P2)}$$

$$S(\theta_1) - c \geq \underline{u}_1(\theta_1) \quad \forall \theta_1 \qquad \text{(interim IR for P1)}$$

Re-arranging the IR constraints gives us upper and lower bounds for c that must be satisfied. The two IR constraints are satisfied iff

$$-s(\theta_1,\theta_2)+S(\theta_1)+\underline{u}_2(\theta_2)\leq c\leq S(\theta_1)-\underline{u}_1(\theta_1)\quad\forall\theta_1,\theta_2$$

In other words, we can find some constant c that satisfies the IR constraints iff:

$$\max_{\theta_1,\theta_2}-s(\theta_1,\theta_2)+S(\theta_1)+\underline{u}_2(\theta_2)\leq\min_{\theta_1}S(\theta_1)-\underline{u}_1(\theta_1).$$

Notice that a lower bound for the left-hand-side is obtained when using $\theta_1=\theta_2=0$, which will also be an upper bound for the right-hand-side. Using our assumption that $\underline{u}_i(0)=0=s(0,0)$, this implies

$$S(0)\leq\max_{\theta_1,\theta_2}-s(\theta_1,\theta_2)+S(\theta_1)+\underline{u}_2(\theta_2)\leq\min_{\theta_1}S(\theta_1)-\underline{u}_1(\theta_1)\leq S(0)$$

Therefore, *if* there exists some constant c that makes the IR constraints hold, it must be that,

$$\max_{\theta_1,\theta_2}-s(\theta_1,\theta_2)+S(\theta_1)+\underline{u}_2(\theta_2)=S(0)=\min_{\theta_1}S(\theta_1)-\underline{u}_1(\theta_1)\tag{Combined IR}$$

Further, if the above equality is satisfied, then only using $c=S(0)$ satisfies the IR constraints. Therefore, the above (Combined IR) equality is necessary and sufficient for there to exist payments that satisfy IR. As a conclusion, sharp upper bounds for the outside options are:

$$\underline{u}_1(\theta_1)\leq S(\theta_1)-S(0)\quad\forall\theta_1\in[0,1]\tag{IIR for delegate}$$
$$\underline{u}_2(\theta_2)\leq\min_{\theta_1}\ s(\theta_1,\theta_2)-S(\theta_1)+S(0)\quad\forall\theta_2\in[0,1]\tag{EPIR for agent}$$

We finally prove that efficiency is achieved for rival but not for non-rival resources when the outside options are zero. Notice that the IIR for the delegate above always holds. Indeed: $u(x,\theta_i)$ increasing implies $s(\theta_1,\theta_2)$ is increasing in both variables and thus $S(\theta_1)>S(0)$. We now check EPIR for the agent. By monotonicity of s we just need to check the left-hand-side at $\theta_2=0$. If resources are rival, by submodularity, for any

$v, \theta_1 \in [0,1]$:

$$s(\theta_1, v) + s(0,0) \leq s(0,v) + s(\theta_1, 0).$$

We now take expectations over v. Using the symmetry of s and the type distributions, and the that $s(0,0) = 0$. we obtain:

$$S(\theta_1) \leq S(0) + s(\theta_1, 0),$$

proving that EPIR holds for the agent. It is clear that if resources are strictly rival the same inequality holds strictly.

Conversely, if the resources are strictly non-rival, strict supermodularity produces, for all $v, \theta_1 \in [0,1]$:

$$s(\theta_1, v) + s(0,0) > s(0,v) + s(\theta_1, 0).$$

By taking expectations we obtain:

$$S(\theta_1) > S(0) + s(\theta_1, 0),$$

showing a failure of EPIR at $\theta_2 = 0$. ■

Proof of Proposition 6

By the bounds in Proposition 5, we must have:

$$r_1\theta_1 = \underline{u}_1(\theta_1) \leq I_F(\theta_1) \quad \forall \theta_1 \in [0,1], \text{ and}$$
$$(1 - r_1)\theta_2 = \underline{u}_2(\theta_2) \leq \min_{\theta_1} \{\max\{\theta_1, \theta_2\} - I_F(\theta_1)\} \quad \forall \theta_2 \in [0,1].$$

If $\theta_1 \leq \theta_2$, the argument in the minimization of the inequality above is decreasing in θ_1, thus $\max\{\theta_1, \theta_2\} - I_F(\theta_1) \geq \theta_2 - I_F(\theta_2)$. If $\theta_2 < \theta_1 < 1$, $\theta_1 - I_F(\theta_1)$ is increasing, because the derivative of $I_F(\theta_1)$ is $F(\theta_1) < 1$, by the assumption of full support. Therefore: $\min_{\theta_1} \{\max\{\theta_1, \theta_2\} - I_F(\theta_1)\} = \theta_2 - I_F(\theta_2)$. Putting the two inequalities above together, we have:

$$I_F(x) = r_1 x,$$

for all $x \in [0,1]$. But that implies F is a binary distribution with support $\{0,1\}$ which is a contradiction. ■

Proof of Proposition 7

By the bounds in Proposition 5, if one can achieve efficiency by delegating to 1, it must be true that:

$$u(r_1,\theta_1) := \underline{u}_1(\theta_1) \leq I_F(\theta_1) \quad \forall \theta_1 \in [0,1], \text{ and} \tag{6}$$

$$u(r_2,\theta_1) := \underline{u}_2(\theta_2) \leq \min_{\theta_1} \{\max\{\theta_1,\theta_2\} - I_F(\theta_1)\} \quad \forall \theta_2 \in [0,1]. \tag{7}$$

By the argument in the proof of Proposition 6, $\min_{\theta_1} \{\max\{\theta_1,\theta_2\} - I_F(\theta_1)\} = \theta_2 - I_F(\theta_2)$. A useful observation is that, by the assumption that $\mu \geq \frac{1}{2}$:

$$x - I_F(x) \geq I_F(x) \quad \forall x \in [0,1] \tag{8}$$

The argument is as follows: both sides of 8 equal 0 at $x = 0$, while at $x = 1$, $1 - I_F(1) = \mu \geq 1 - \mu = I_F(1)$. Because $x - I_F(x)$ is concave and increasing, while $I_F(x)$ is convex and increasing, the two cannot cross at any $x \in (0,1)$.

Suppose we instead delegate to player 2 and let player 1 be the agent. Then, player 1's ex-post payoff satisfies

$$u(r_1,v) \leq I_F(v) \leq v - I_F(v) \; \forall v \in [0,1],$$

where the first inequality comes from 6 and the second inequality from the observation 8. Therefore, we do satisfy ex-post individual rationality for player 1. For player 2:

$$u(r_2,v) \leq u(r_1,v) \leq I_F(v),$$

where the first inequality comes from $r_2 \leq r_1$, and the second inequality follows from 6. ■