Robustness in Costly Contracting

Francisco del Villar

June 19, 2024

Link to the latest version

Abstract

A principal can contract with one of several agents based on observed output. If she did not hire anyone, she would take a baseline action that is also available to agents and produces some output. Agents know what happens under the baseline action, but the principal does not. In addition, the principal ignores what other actions are available to agents. In this situation, simply giving a contract away for free can make the principal worse off. To avoid this, she auctions off the contract. She cannot pinpoint what agents know about their contract values and evaluates contract and auction pairs according to her payoff in worst-case scenarios. We find that a first-price auction of a full-benefit contract, which pays all of the realized output, is optimal with a payoff guarantee of zero. Any other contract can make the principal worse off, no matter how she sells it. These findings guide those who want to outsource efforts towards improving observable outcomes but face nontrivial opportunity costs of contracting and uncertainty that is hard to quantify.

Consider a principal who thinks about hiring an agent to produce output. Hiring may or may not be a good idea, depending on what the principal would do in the baseline — if she did not hire anyone. The principal's baseline action is typically assumed to be costless and produce the lowest possible output, zero (Carroll 2015). Here, contracting is a safe business for her. In particular, any output-sharing agreement is surely beneficial. The agent might produce something, in which case the principal is better off, or he might produce nothing, in which case the principal is as well off as if she did not hire anyone.

An opposite conclusion arises if the principal's baseline action produces some output. Here, any contract can make her worse off. We can see this in the following example. Say the principal leads an environmental organization and wishes to hire an agent to preserve a section of rainforest spanning 100 acres. If she did not hire an agent, she would turn to other matters, and some rainforest — say 40 acres — would still survive by the end of the year. Now consider any contract that pays according to the preserved acreage. The principal ignores what the agent can and cannot do to preserve, and it might be optimal for a contracted agent to follow the principal's baseline action, do nothing, and get paid for the 40 preserved acres. This situation harms the principal because she has to pay for an outcome that would have happened anyway.

To complicate matters, our principal is uncertain about the preserved acreage under the baseline action. How could she hire an agent and never be worse off? The key is to have agents somehow reimburse her for the contract payment under the baseline action. The principal's baseline action is available to all agents. Hence, every agent is willing to pay at least the contract payment under the baseline action to own the contract. Our principal exploits this insight by having agents compete for the contract in an auction.

In our model, the principal knows that the baseline action is available to agents but ignores what other actions are also available. In addition, she ignores what agents know about each others' contract values and cannot select a bidding equilibrium. To select a contract and an auction in this environment, she adopts a maxmin criterion and evaluates contract and auction pairs according to their worst-possible payoff.

Our main finding is that it is optimal for the principal to sell a full-benefit contract, which pays all of the realized output, with a first-price auction. This contract and auction pair gives the principal a payoff guarantee of zero. The result holds across various extensions of the model in which we allow the principal to have more knowledge about the baseline action and about the set of actions available to agents. The full-benefit contract is uniquely optimal: any other contract can harm the principal, no matter how she sells it. However, we do not show that the first-price auction is uniquely optimal. In principle, any auction whose revenue exceeds the lowest possible valuation of agents across information structures consistent with a common prior and equilibria also produces a payoff guarantee of zero when paired with the full-benefit contract. The first-price auction satisfies this property by virtue of its revenue guarantee, characterized by Bergemann, Brooks, and Morris (2017).

The intuition for our main finding is simple. Under the full-benefit contract, the principal surrenders her entire benefit from the realized outcome to the agent. Her payoff then boils down to the auction's revenue net of her opportunity cost of allocating the contract, which is her payoff under the baseline action. Agents' valuations for the full-benefit contract are bounded below by their payoff under the baseline action. But this payoff exactly equals the principal's opportunity cost of allocating the contract. With a first-price auction, revenue weakly exceeds this quantity, so the principal is weakly better off.

The contribution of this paper is twofold. First, our main finding guides parties who wish to outsource efforts towards improving an observable outcome but face nontrivial opportunity costs of contracting and uncertainty that is hard to quantify. Second, our model provides a framework to simultaneously address designer uncertainty to players' actions and uncertainty to players' information in a non-bayesian fashion. We define type spaces to include both a common prior information structure that specifies a hierarchy of agent beliefs about their values of the contract as well as a specification of the counterfactual outcomes that agents would produce under contract and that the principal would produce if she did not hire anyone. Formally, our principal considers a class of contract-specific type spaces and judges contract and auction pairs according to her payoff in the worst possible type space.

This paper joins the literature on robust mechanism design, where the designer evaluates alternative mechanisms according to their worst-case performance (see Carroll 2019 for a review). We depart from the contracting setting of Carroll (2015) by allowing the principal's baseline action to produce output beyond its lowest possible level, normalized to zero. The literature that studies properties of auctions across classes of common prior information structures is closely related to this paper (Bergemann, Brooks, and Morris 2020, Brooks and Du 2021). In particular, Bergemann, Brooks, and Morris (2017) is our stepping stone to establish the optimality of the first-price auction. Another related literature studies the design and implementation of Advance Market Commitments — contracts that pay a fixed price to anyone who delivers a unit of a good or service that may not exist yet (Kremer and Glennerster 2004, Kremer, Levin, and Snyder 2020, Kremer, Levin, and Snyder 2022).

We present the basic model in section 1. Section 2 presents our main results, and section 3 extends the model in a number of directions. Section 4 concludes.

1 Model

We consider a principal who is risk-neutral and values an outcome that belongs to the interval $[0, \overline{y}]$, where $\overline{y} > 0$. We can interpret the outcome as output or as the principal's monetary benefit from an observable outcome. $n \geq 2$ agents are available to improve the outcome, and the principal's problem is to hire and motivate one agent. She does so by issuing a contract and allocating it to one agent through an auction. An agent who values the contract at $v \in \mathbb{R}$ has preferences over probabilities of receiving the contract $q_i \in [0, 1]$ and auction transfers $t_i \in \mathbb{R}$ given by $vq_i - t_i$.

1.1 Contracts and Auctions

A contract is a continuous function $w : [0, \overline{y}] \mapsto [0, \infty)$ that requires the principal to pay the contract holder $w(y) \geq 0$ dollars if the realized outcome is y. An auction is a tuple $a = (\mathcal{M}, q, t)$, where $\mathcal{M}$ is a set of messages available to agents, $q : \mathcal{M}^n \mapsto [0, 1]^n$ is the allocation rule, and $t : \mathcal{M}^n \mapsto \mathbb{R}^n$ is the transfer rule.[1] Given a profile of messages $m = (m_1, \dots, m_n) \in \mathcal{M}^n$, $q_i(m)$ is the probability that $i \in \{1, \dots, n\}$ gets the contract and $t_i(m)$ is the dollar amount that i owes the principal.

Our auctions allocate the good with certainty, so $\sum_{i=1}^n q_i(m) = 1$ for all profiles of messages m. We consider auctions that may not allocate the contract in section 3. Moreover, agents can send a message $0 \in \mathcal{M}$ that ensures a weakly negative transfer: $t_i(0, m_{-i}) \leq 0$ for all $i \in \{1, \dots, n\}$ and for all messages of agents other than i, $m_{-i} \in \mathcal{M}^{n-1}$. Agents' contract valuations will be non-negative, so this condition implies that our auctions satisfy participation security (Brooks and Du 2021).

1.2 Contract Valuations

We summarize what agents know about their values of a contract with a Harsanyi type space, which includes a set $\mathcal{S}$ of signals, a joint distribution over profiles of signal $\psi \in \Delta(\mathcal{S}^n)$ and a value function $\widehat{v}_i : \mathcal{S}^n \mapsto \mathbb{R}$ for each agent $i \in \{1, \dots, n\}$. $\widehat{v}_i(s)$ is the expected value of the contract to agent i, conditional on the profile of signals $s = (s_1, \dots, s_n) \in \mathcal{S}^n$.

[1] I have to endow set $\mathcal{M}$ with a σ-algebra so I can define strategies later on and take integrals. Bergemann, Brooks, and Morris (2017) do it like this (copy-pasted): All sets considered in this paper are regarded as topological spaces with their standard topologies, wherever applicable, and endowed with the Borel σ-algebra. For a topological space $\mathcal{X}$, $\Delta(\mathcal{X})$ denotes the set of Borel probability measures on $\mathcal{X}$, endowed with the weak-* topology. For a measure $\mu \in \Delta(\mathcal{X})$ and a measurable function $f : \mathcal{X} \mapsto \mathbb{R}$, we denote the integral of f with respect to μ by $\int_{\mathcal{X}} f(x)\mu(dx)$.

1.3 Actions of Agents and of the Principal

We do not explicitly discuss the decisions of agents under contract. Instead, we specify what agents would accomplish under contract and what the principal would accomplish if she did not hire anyone.

Conditional on signal profile $s \in \mathcal{S}^n$, agent i would produce an expected outcome $\widehat{y}_i(s) \in [0, \overline{y}]$ if she received the contract and would earn an expected pay equal to $\widehat{w}_i(s) \geq 0$. If the principal did not hire an agent, she would take a prespecified baseline action. Conditional on a signal profile $s \in \mathcal{S}^n$, the baseline action produces an expected outcome $\widehat{y}_b(s) \in [0, \overline{y}]$ at an expected cost of $\widehat{c}_b(s) \geq 0$. This action is also available to agents. An agent who takes the baseline action earns an expected contract payment equal to $\widehat{w}_b(s) \geq 0$, conditional on $s \in \mathcal{S}^n$.

1.4 Type Spaces

We can now complete our definition of a type space: it is a tuple $T = (\mathcal{S}, \psi, \widehat{v}, \widehat{y}, \widehat{w}, \widehat{c}_b)$, where $\widehat{v} = (\widehat{v}_1, \ldots, \widehat{v}_n)$, $\widehat{y} = (\widehat{y}_b, \widehat{y}_1, \ldots, \widehat{y}_n)$ and $\widehat{w} = (\widehat{w}_b, \widehat{w}_1, \ldots, \widehat{w}_n)$.

A type space T is consistent with contract w if for every profile of signals $s \in \mathcal{S}^n$, there are probability distributions $F_b, F_1, \ldots, F_n \in \Delta([0, \overline{y}])$ such that

$$\begin{aligned} \widehat{y}_b(s) &= \int_0^{\overline{y}} y F_b(dy), \quad \widehat{w}_b(s) = \int_0^{\overline{y}} w(y) F_b(dy) \quad \text{and} \\ \widehat{y}_i(s) &= \int_0^{\overline{y}} y F_i(dy), \quad \widehat{w}_i(s) = \int_0^{\overline{y}} w(y) F_i(dy) \quad \text{for all } i \in \{1, \ldots, n\}, \end{aligned} \tag{1}$$

that is, if expected outcomes and contract payments belong to the convex hull of the graph of contract w.

Our type spaces satisfy three conditions. First, it is costly for agents to act, so that

$$\widehat{w}_i(s) \geq \widehat{v}_i(s) \text{ for all } i \text{ and } s \in \mathcal{S}^n. \tag{2}$$

Second, agents can follow the principal's baseline action and may also take a costless action. Since contracts are non-negative, this means that

$$\widehat{v}_i(s) \geq \max\{0,\ \widehat{w}_b(s) - \widehat{c}_b(s)\} \text{ for all } i \text{ and } s \in \mathcal{S}^n. \tag{3}$$

Finally, the baseline action is beneficial to the principal, meaning that

$$\int_{\mathcal{S}^n} [\widehat{y}_b(s) - \widehat{c}_b(s)] \psi(ds) \geq 0. \tag{4}$$

The set of admissible type spaces for contract w is then

$$\mathcal{T}_w = \{T : T \text{ satisfies (1), (2), (3) and (4)}\}.$$

The following property of type spaces will be central to our analysis. A type space T features *common baseline knowledge* if the consequences of the baseline action are common knowledge among agents:[2]

$$\widehat{y}_b(s) = \widehat{y}_b(s'),\ \widehat{c}_b(s) = \widehat{c}_b(s') \text{ and } \widehat{w}_b(s) = \widehat{w}_b(s') \text{ for all } s, s' \in \mathcal{S}^n. \tag{5}$$

1.5 Bidding Behavior

An auction a and a type space T specify a game of incomplete information among agents. A strategy for agent i is a mapping $\sigma_i : \mathcal{S} \mapsto \Delta(\mathcal{M})$ that associates a probability distribution over messages to each signal observation. Given a profile of strategies $\sigma = (\sigma_1, \ldots, \sigma_n)$, agent i's ex-ante expected utility is

$$u_i(\sigma \,|\, T, a) = \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} [q_i(m)\widehat{v}_i(s) - t_i(m)]\sigma_1(dm_1|s_1) \cdots \sigma_n(dm_n|s_n)\psi(ds).$$

A Bayes Nash equilibrium is a profile of strategies σ such that $u_i(\sigma) \geq u_i(\sigma_i', \sigma_{-i})$ for every agent i and strategy σ_i', where σ_{-i} is the profile of strategies of agents other than i. The set of Bayes Nash equilibria is $\Sigma(T, a)$.

1.6 Timing

First, the principal selects a contract and an auction format and announces them. Agents privately observe their signal realizations and participate in the auction. After they bid, the contract is allocated, and the auction transfers are determined. Then, the contracted agent acts. Finally, the outcome occurs, and all transactions take place. Agents pay the principal according to their auction messages and the transfer rule, and the principal pays the contracted agent according to the realized outcome and the contract.

In this sequence of events, agents delay their auction transfers to the principal until after the outcome occurs instead of paying immediately after the auction. This choice of timing allows us to ignore individual attitudes towards payments made at the end of the auction and those made after the outcome occurs.

[2] In general, baseline outcomes, costs, and contract payments could depend on a signal component that is common knowledge among agents. Here, we proceed as if agents have already observed the realization of this common signal.

1.7 Objective of the Principal

For a given auction a, the principal's expected payoff at type space T and profile of strategies σ is $\Pi(T, \sigma \,|\, a)$, given by

$$\begin{aligned}\Pi(T,\sigma\,|\,a) &= \int\limits_{\mathcal{S}^n}\int\limits_{\mathcal{M}^n}\sum_{i=1}^{n}\Big(q_i(m)\big[\widehat{y}_i(s)-\widehat{w}_i(s)\big]+t_i(m)\Big)\sigma_1(dm_1|s_1)\cdots\sigma_n(dm_n|s_n)\psi(ds)\\ &\qquad -\int\limits_{\mathcal{S}^n}\big[\widehat{y}_b(s)-\widehat{c}_b(s)\big]\psi(ds)\\ &= \int\limits_{\mathcal{S}^n}\int\limits_{\mathcal{M}^n}\sum_{i=1}^{n}\Big(q_i(m)\big[\widehat{y}_i(s)-\widehat{w}_i(s)-(\widehat{y}_b(s)-\widehat{c}_b(s))\big]+t_i(m)\Big)\sigma(dm|s)\psi(ds), \quad (6)\end{aligned}$$

where $\sigma(dm|s) = \sigma_1(dm_1|s_1)\cdots\sigma_n(dm_n|s_n)$.

The principal is uncertain about both the type space and the strategies that agents would follow at any given auction and type space. Given contract w, she knows that the type space belongs to the subset of admissible type spaces that feature common baseline knowledge

$$\mathcal{T}_w^* = \{T \in \mathcal{T}_w :\ T \text{ satisfies (5)}\}.$$

She also knows that agents follow Bayes Nash equilibrium strategies at any given auction and type space.

That said, she does not rely on a prior on the set of possible type spaces and equilibria, $\{(T,\sigma) :\ T \in \mathcal{T}_w^*,\ \sigma \in \Sigma(T,a)\}$, to evaluate contract w and auction a. Instead, she evaluates a contract w and an auction a with the *payoff guarantee* $\underline{\Pi}(w,a)$, which is the worst-possible expected payoff she could obtain by looking across possible type spaces and equilibria:

$$\underline{\Pi}(w,a) = \inf\Big\{\Pi(T,\sigma\,|\,a) :\ T \in \mathcal{T}_w^* \text{ and } \sigma \in \Sigma(T,a)\Big\}$$

if the infimum exists. Otherwise, $\underline{\Pi}(w,a) \equiv -\infty$.

We now turn to our central task, which is to find contracts and auctions that maximize the principal's payoff guarantee.

2 Analysis

This section discusses our main result, which is that a first-price auction of a full-benefit contract is optimal for a principal who knows that agents have common baseline knowledge. In the first-price auction, messages are non-negative bids, $\mathcal{M} = [0, \infty)$; the principal allocates the contract among those who bid the most, $m_i < \max\{m_1, \ldots, m_n\}$ implies that $q_i(m) = 0$; and her revenue equals the highest bid, $t_1(m) + \cdots + t_n(m) = \max\{m_1, \ldots, m_n\}$. In turn, a full-benefit contract w satisfies $w(y) = y + k$, $k \geq 0$, for all $y \in [0, \overline{y}]$.

Theorem 1. *It is optimal for the principal to sell any full-benefit contract with a first-price auction.*

To prove this result, we first argue that the principal can be weakly worse off from any contract and auction pair due to the possibility that agents optimally want to mimic the principal's baseline action. Then, we argue that the principal must be weakly better off from a first-price auction of a full-benefit contract. The first-price auction emerges because its revenue guarantee, characterized by Bergemann, Brooks, and Morris (2017), exceeds the lowest possible valuation consistent with a common prior. The full-benefit contract $w(y) = y$ is special because it reduces the principal's payoff to the auction's revenue net of her opportunity cost of contracting, which is her net benefit of the baseline action. When agents have common baseline knowledge, they know that their values of the full-benefit contract exceed this opportunity cost. So, if the principal sells such a contract with a first-price auction, she cannot be worse off. When the contract also includes a lump sum transfer, so that $w(y) = y + k$ with $k > 0$, the argument is virtually unchanged: agents' contract valuations increase by k, and they transfer k back to the principal when they bid in the first-price auction.

Lemma 1. *For every contract w, there exists a type space $T \in \mathcal{T}_w^*$ such that $\Pi(T, \sigma \mid a) \leq 0$ for every auction a and equilibrium $\sigma \in \Sigma(T, a)$.*

Proof. Fix $y_b \in [0, \overline{y}]$, $w_b = w(y_b)$ and c_b satisfying $0 \leq c_b \leq \min\{y_b, w_b\}$. At type space T, agents mimic the baseline action: for every agent $i \in \{1, \ldots, n\}$ and signal profile $s \in \mathcal{S}^n$,

$$\widehat{y}_i(s) = \widehat{y}_b(s) = y_b, \quad \widehat{w}_i(s) = \widehat{w}_b(s) = w_b, \quad \widehat{c}_b(s) = c_b, \quad \text{and} \quad \widehat{v}_i(s) = w_b - c_b.$$

We can see that T satisfies conditions (1) through (5), so $T \in \mathcal{T}_w^*$. T is bad for the principal because it implies a zero-sum game between her and the agents: for any auction a

and every profile of strategies σ,

$$\begin{aligned}\Pi(T,\sigma\,|\,a) &= \int_{\mathcal{S}^n}\int_{\mathbb{R}^n_+}\sum_{i=1}^n\Big[q_i(m)\Big(\widehat{y}_i(s)-\widehat{w}_i(s)-(\widehat{y}_b(s)-\widehat{c}_b(s))\Big)+t_i(m)\Big]\sigma(dm|s)\psi(ds)\\ &= \int_{\mathcal{S}^n}\int_{\mathbb{R}^n_+}\sum_{i=1}^n -\Big(q_i(m)\widehat{v}(s)-t_i(m)\Big)\sigma(dm|s)\psi(ds)\\ &= -\sum_{i=1}^n u_i(\sigma\,|\,T,a).\end{aligned}$$

Agents may always send message $0\in\mathcal{M}$ and obtain a non-negative payoff. In an equilibrium $\sigma\in\Sigma(T,a)$, we conclude that $\Pi(T,\sigma\,|\,a)\le 0$. ■

Lemma 1 shows that the payoff guarantee of any contract and auction pair is weakly negative. We shall now see that the principal must be weakly better off under certain auctions of a full-benefit contract.

We say that an auction $a=(\mathcal{M},q,t)$ has a *value-guaranteed revenue* if agents' valuations bound equilibrium revenues: for every contract w and type space $T\in\mathcal{T}_w$,

$$\begin{cases}\sigma\in\Sigma(T,a)\quad\text{and}\\ \widehat{v}_i(s)\ge x \text{ for all } i \text{ and } s\in\mathcal{S}^n\end{cases}\qquad\text{implies}\quad \int_{\mathcal{S}^n}\int_{\mathcal{M}^n}\Big(\sum_{i=1}^n t_i(m)\Big)\sigma(dm|s)\psi(ds)\ge x.$$

Our next result is that the principal cannot be worse off if she sells a full-benefit contract with an auction that has a value-guaranteed revenue.

Lemma 2. *Consider a contract $w(y)=y+k$ with $k\ge 0$ and let auction a have a value-guaranteed revenue. For every type space $T\in\mathcal{T}^*_w$ and equilibrium $\sigma\in\Sigma(T,a)$, $\Pi(T,\sigma\,|\,a)\ge 0$.*

Proof. Fix a type space $T\in\mathcal{T}^*_w$. T satisfies (5), so there are $y_b\in[0,\overline{y}]$ and $c_b\ge 0$ such that $\widehat{y}_b(s)=y_b$ and $\widehat{c}_b(s)=c_b$ for all $s\in\mathcal{S}^n$. Moreover, T is consistent with contract w so that $\widehat{w}_i(s)=\widehat{y}_i(s)+k$ and $\widehat{w}_b(s)=y_b+k$ for all $i\in\{1,\dots,n\}$ and $s\in\mathcal{S}^n$. For any profile of strategies σ, it follows that

$$\begin{aligned}\Pi(T,\sigma\,|\,w,a) &= \int_{\mathcal{S}^n}\int_{\mathcal{M}^n}\sum_{i=1}^n\Big(q_i(m)\big[\widehat{y}_i(s)-\widehat{w}_i(s)-(\widehat{y}_b(s)-\widehat{c}_b(s))\big]+t_i(m)\Big)\sigma(dm|s)\psi(ds)\\ &= \int_{\mathcal{S}^n}\int_{\mathcal{M}^n}\Big(\sum_{i=1}^n t_i(m)\Big)\sigma(dm|s)\psi(ds)-(y_b+k-c_b).\end{aligned}$$

Now, T satisfies (3), meaning that $\widehat{v}_i(s) \geq \widehat{w}_b(s) - \widehat{c}_b(s) = y_b - c_b + k$ for all i and $s \in \mathcal{S}^n$. Since auction a has a value-guaranteed revenue, we conclude that $\Pi(T, \sigma \,|\, a) \geq 0$ for every equilibrium σ. ∎

Taken together, Lemmas 1 and 2 imply that selling a full-benefit contract with an auction that has a value-guaranteed revenue yields the best payoff guarantee, which is zero. But is there an auction that satisfies this property?

The second-price (Vickrey) auction does not. If agents' values of the contract are private, it is a weakly dominant strategy for them to bid their value.[3] If agents followed this strategy, the auction's revenue would naturally exceed any lower bound on agents' valuations. Unfortunately, this conclusion is fragile. As is well-known, even with private values, the second price auction features bidding ring equilibria where one agent bids an astronomical amount, others bid below any possible value they might hold, and revenue is below any possible valuation. We can also find low revenues if we look past private values. For example, when agents have a common value of the contract and one agent is better informed about the value than others (Engelbrecht-Wiggans, Milgrom, and Weber 1983, Bergemann, Brooks, and Morris 2019).

Bergemann, Brooks, and Morris (2017) showed that, in contrast with the second-price auction, revenue in the first-price auction cannot be too low. We can use their results to show that it has a value-guaranteed revenue.

Lemma 3. *The first-price auction has a value-guaranteed revenue.*

Proof. Fix a contract w and consider a type space $T \in \mathcal{T}_w$. Let $x \in \mathbb{R}$ be such that $\widehat{v}_i(s) \geq x$ for all i and $s \in \mathcal{S}^n$ and define $\underline{v} \equiv \max\{0, x\}$. T satisfies (3) so that values are non-negative and $\widehat{v}_i(s) \geq \underline{v}$ for all i and $s \in \mathcal{S}^n$. w is a continuous function, so we can set $\overline{v} = \max_{y \in [0,\overline{y}]} w(y)$. It follows that $\widehat{v}_i(s) \in [\underline{v}, \overline{v}] \subseteq [0, \infty)$ for all i and s, so agents' valuations are consistent with a common prior $\mu \in \Delta([\underline{v}, \overline{v}]^n)$.

By definition, the first-price auction's equilibrium revenue at type space T must exceed its revenue guarantee $\underline{R}$ at the common prior μ, which was shown by Bergemann, Brooks, and Morris (2017) to be:

$$\underline{R} = \int_{\underline{v}}^{\overline{v}} \underline{\beta}(x) Q(dx),$$

where Q is the cumulative distribution function of the average of the $n-1$ lowest values implied by μ and $\underline{\beta}$ is the minimum winning bid function. Specifically, $Q : [\underline{v}, \overline{v}] \mapsto [0, 1]$ is

[3] Formally, type space T features private values if for every agent i and profile of signals $s = (s_1, \ldots, s_n) \in \mathcal{S}^n$, $\widehat{v}_i(s) = \widehat{v}_i(s_i)$. That is, the information of others does not influence an agent's contract value.

given by

$$Q(x) = \mu\big(\{v \in [\underline{v}, \overline{v}]^n : \alpha(v) \leq x\}\big),$$

where

$$\alpha(v) = \frac{1}{n-1}\left(\sum_{i=1}^{n} v_i - \max\{v_1, \dots, v_n\}\right).$$

The minimum winning bid function $\underline{\beta}$ satisfies

$$\underline{\beta}(z) = \int_{\underline{v}}^{z} x \frac{Q^{(n-1)/n}(dx)}{Q^{(n-1)/n}(z)}$$

for all $z \in [\underline{v}, \overline{v}]$. Given any such z, $F_z : [\underline{v}, z] \mapsto [0,1]$ given by $F_z(x) \equiv Q^{(n-1)/n}(x)/Q^{(n-1)/n}(z)$ is a cumulative distribution function. Therefore, $\underline{\beta}(z)$ is the expectation of a random variable distributed according to F_z on the interval $[\underline{v}, z]$. It follows that $\underline{\beta}(z) \geq \underline{v}$ and we conclude that for any equilibrium $\sigma \in \Sigma(T, a)$,

$$\int_{\mathcal{S}^n}\int_{\mathcal{M}^n} \left(\sum_{i=1}^{n} t_i(m)\right)\sigma(dm|s)\psi(ds) \;\geq\; \underline{R} \;=\; \int_{\underline{v}}^{\overline{v}} \underline{\beta}(x)Q(dx) \;\geq\; \int_{\underline{v}}^{\overline{v}} \underline{v}Q(dx) \;=\; \underline{v} \;\geq\; x.$$

■

Lemmas 1, 2 and 3 prove Theorem 1. The first-price auction is optimal in that it has a value-guaranteed revenue. It is not uniquely optimal because the principal could achieve the best payoff guarantee by selling a full-benefit contract with any other auction with a value-guaranteed revenue.

We now turn our attention to the contract. Could other contracts also be optimal with an appropriate auction pairing? The answer is negative.

Proposition 1. *Consider a contract w such that $w(y_1) - w(y_0) \neq y_1 - y_0$ for some $y_0, y_1 \in [0, \overline{y}]$. For every auction a, there exists a type space $T \in \mathcal{T}_w^*$ such that $\Pi(T, \sigma \,|\, a) < 0$ for every equilibrium $\sigma \in \Sigma(T, a)$.*

Proof. Take $y_0 < y_1$ and consider two cases.

Case 1 Suppose that $w(y_1) - w(y_0) < y_1 - y_0$. The type space we construct is such that agents prefer to produce the smaller outcome y_0, whereas the baseline action would produce outcome y_1. Moreover, the principal is better off under the baseline action even if she did not have to make a contract payment for the smaller outcome y_0. This means that contracting makes her worse off even if she can fully extract agents' values.

That is, $\widehat{y}_i(s) = y_0$, $\widehat{w}_i(s) = w(y_0) = \widehat{v}_i(s)$, $\widehat{y}_b(s) = y_1$ and $\widehat{w}_b(s) = w(y_1)$ for all $i \in \{1, \ldots, n\}$ and $s \in \mathcal{S}^n$. Moreover, $\widehat{c}_b(s) = c \in \big[\max\{0,\, w(y_1) - w(y_0)\},\; y_1 - y_0\big)$ for all $s \in \mathcal{S}^n$. Clearly, T satisfies (1) through (5), so that $T \in \mathcal{T}_w^*$. Now fix an auction a and an equilibrium $\sigma \in \Sigma(T, a)$. It follows that

$$\begin{aligned}
\Pi(T, \sigma \,|\, a) &= \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} \Big(q_i(m) \big[\widehat{y}_i(s) - \widehat{w}_i(s) - (\widehat{y}_b(s) - \widehat{c}_b(s))\big] + t_i(m) \Big) \sigma(dm|s)\psi(ds) \\
&= y_0 - w(y_0) - (y_1 - c) + \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} t_i(m)\sigma(dm|s)\psi(ds) \\
&\leq y_0 - w(y_0) - (y_1 - c) + \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} q_i(m)\widehat{v}_i(s)\sigma(dm|s)\psi(ds) \\
&= (y_1 - c) - y_0 \\
&< 0,
\end{aligned}$$

where the weak inequality follows from the fact that agents have a non-negative surplus in equilibrium.

Case 2 Now suppose that $y_1 - y_0 < w(y_1) - w(y_0)$. Here, we construct a type space such that agents produce the larger outcome y_1 and the baseline action produces y_0. Agents' costs of producing y_1 drive down their contract values to the point where any auction's revenue cannot compensate the principal for the relatively large contract payment, making her worse off.

Let T be such that $\widehat{y}_i(s) = y_1$, $\widehat{w}_i(s) = w(y_1)$, $\widehat{y}_b(s) = y_0$ and $\widehat{w}_b(s) = w(y_0)$ for all $i \in \{1, \ldots, n\}$ and $s \in \mathcal{S}^n$. The baseline action is costless, so $\widehat{c}_b(s) = 0$ for all s. Moreover, agents' values satisfy $\widehat{v}_i(s) = w(y_1) - c$ for all i and s, where $c \in (y_1 - y_0, w(y_1) - w(y_0)]$. Once again, T satisfies (1) through (5), so that $T \in \mathcal{T}_w^*$. For any auction a and equilibrium

$\sigma \in \Sigma(T, a)$, we have that

$$\begin{aligned}
\Pi(T, \sigma \mid a) &= \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} \Big(q_i(m) \big[\widehat{y}_i(s) - \widehat{w}_i(s) - (\widehat{y}_b(s) - \widehat{c}_b(s)) \big] + t_i(m) \Big) \sigma(dm|s) \psi(ds) \\
&= y_1 - w(y_1) - y_0 + \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} t_i(m) \sigma(dm|s) \psi(ds) \\
&\leq y_1 - w(y_1) - y_0 + \int\limits_{\mathcal{S}^n} \int\limits_{\mathcal{M}^n} \sum_{i=1}^{n} q_i(m) \widehat{v}_i(s) \sigma(dm|s) \psi(ds) \\
&= [y_1 - w(y_1)] + [w(y_1) - c] - y_0 \\
&< 0.
\end{aligned}$$

The weak inequality follows again from the fact that agents have a non-negative surplus in equilibrium. ■

Up to now, we have assumed that type spaces feature common baseline knowledge, meaning that the consequences of the baseline action are common knowledge among agents. If we drop this assumption, it turns out that any contract and auction pair can make the principal worse off.

Proposition 2. *For every contract w, there exists a type space $T \in \mathcal{T}_w$ such that $\Pi(T, \sigma \mid a) < 0$ for any auction a and equilibrium $\sigma \in \Sigma(T, a)$.*

Proof. By Proposition 1, we need only consider a full-benefit contract satisfying $w(y) = y + k$ for all $y \in [0, \overline{y}]$, with $k \geq 0$. In our type space T, agents mimic the principal's baseline action, which is costless. That is, $\widehat{y}_i(s) = \widehat{y}_b(s)$, $\widehat{c}_b(s) = 0$, $\widehat{w}_i(s) = \widehat{w}_b(s) = \widehat{y}_b(s) + k$ and $\widehat{v}_i(s) = \widehat{v}(s) = \widehat{y}_b(s) + k$ for every $i \in \{1, \ldots, n\}$ and $s \in \mathcal{S}^n$. T satisfies (1) through (4), so that $T \in \mathcal{T}_w$. At T, agents have a common value of contract w. Furthermore, T implies a zero-sum game between the principal and the agents: for any auction a and every profile of strategies σ,

$$\begin{aligned}
\Pi(T, \sigma \mid a) &= \int\limits_{\mathcal{S}^n} \int\limits_{\mathbb{R}_+^n} \sum_{i=1}^{n} \Big[q_i(m) \Big(\widehat{y}_i(s) - \widehat{w}_i(s) - (\widehat{y}_b(s) - \widehat{c}_b(s)) \Big) + t_i(m) \Big] \sigma(dm|s) \psi(ds) \\
&= \int\limits_{\mathcal{S}^n} \int\limits_{\mathbb{R}_+^n} \sum_{i=1}^{n} - \Big(q_i(m) \widehat{v}(s) - t_i(m) \Big) \sigma(dm|s) \psi(ds) \\
&= - \sum_{i=1}^{n} u_i(\sigma \mid T, a).
\end{aligned}$$

We will further specify T so that agents obtain a strictly positive surplus in every auction a and equilibrium $\sigma \in \Sigma(T, a)$.

Consider the maximum signal model of Bulow and Klemperer (2002). The highest signal reveals the value of the contract: $\widehat{v}(s_1, \ldots, s_n) = \max\{s_1, \ldots, s_n\}$ and signals are independent draws from a uniform distribution with support equal to $[k,\, \overline{y}+k]$. In this model, Bergemann, Brooks, and Morris (2020) show (Theorem 3) that an upper bound on revenue across auctions and equilibria is

$$\overline{R} \equiv \int_{k}^{\overline{y}+k} \phi(v) g(v) dv,$$

where

$$\phi(v) = v - \int_{v}^{\overline{y}+k} \frac{\overline{y} - (x-k)}{x-k} dx$$

is interpreted as a virtual value and g is the probability density function of the maximum signal, $g(v) = n(v-k)^{n-1}/\overline{y}^n$. In this model, the principal is unable to extract the entire bidders' surplus, regardless of the auction format and equilibrium thereof:

$$\int_{k}^{\overline{y}+k} v g(v) dv - \overline{R} = \int_{k}^{\overline{y}+k} \int_{v}^{\overline{y}+k} \frac{\overline{y} - (x-k)}{x-k} dx g(v) dv > 0.$$

Therefore, $\sum_{i=1}^{n} u_i(\sigma \,|\, T, a) > 0$ and $\Pi(T, \sigma \,|\, a) < 0$ for every auction a and equilibrium $\sigma \in \Sigma(T, a)$. ■

3 Extensions

Section 2 discusses the contract and auction design problem of a principal who is quite ignorant. She knows that agents can also take the baseline action and that its consequences are common knowledge among them. However, she ignores these consequences; she ignores what other actions are available to them; and she ignores what they know about each others' contract values. Moreover, section 2 restricts attention to auctions that allocate the contract with certainty.

In this section, we ask what happens when the principal can also choose auctions that may not allocate the contract; when she knows more about her payoff under the baseline action; and when she knows about other actions that are available to agents.

3.1 Can-Keep Auctions

Our model excludes auctions that may not allocate the contract, like auctions where the principal sets a reserve price under which she keeps the contract and takes the baseline action. In a can-keep auction, the allocation rule q satisfies

$$\sum_{i=1}^{n} q_i(m) \leq 1$$

for all message profiles $m \in \mathcal{M}^n$. Note that, given a can-keep auction a the principal's expected payoff at type space T and strategies σ is still given by (6):

$$\Pi(T, \sigma \,|\, a) = \int_{\mathcal{S}^n} \int_{\mathcal{M}^n} \sum_{i=1}^{n} \Big(q_i(m) \big[\widehat{y}_i(s) - \widehat{w}_i(s) - (\widehat{y}_b(s) - \widehat{c}_b(s)) \big] + t_i(m) \Big) \sigma(dm|s) \psi(ds).$$

Could the principal benefit from this added flexibility? Not really: the argument in Lemma 1, which establishes that the principal can be weakly worse off from any contract and auction pair, does not rely on auctions allocating the contract with certainty. We record this as a corollary.

Corollary 1. *For every contract w, there exists a type space $T \in \mathcal{T}_w^*$ such that $\Pi(T, \sigma \,|\, a) \leq 0$ for every can-keep auction a and equilibrium $\sigma \in \Sigma(T, a)$.*

Proof. Identical to that of Lemma 1. ■

This result implies that the payoff guarantee from any contract and auction pair is bounded above by zero so that it continues to be optimal for the principal to sell a full-benefit contract with a first-price auction.

3.2 Known Baseline Action

In this section, we allow the principal to know the consequences of the baseline action. This knowledge enlarges the set of optimal contracts and auctions. It affords the principal flexibility to sell any output-sharing agreement through an auction with a reserve price set to her payoff under the baseline action, i.e., her opportunity cost of contracting. However, it does not allow the principal to improve her payoff guarantee: optimal contract and auction pairs continue to have a payoff guarantee of zero. This is true because it may still be optimal for agents to mimic the baseline action or otherwise produce small outcomes, meaning that a version of Lemma 1 still applies. In particular, a first-price auction of a full-benefit contract continues to be optimal.

Formally, an action involves a probability distribution $F_b \in \Delta([0, \overline{y}])$ and a cost $c_b \geq 0$ (Carroll 2015). A type space T is consistent with (F_b, c_b) if

$$\begin{aligned} \int_{\mathcal{S}^n} \widehat{y}_b(s)\psi(ds) &= \int_0^{\overline{y}} y F_b(dy), \\ \int_{\mathcal{S}^n} \widehat{w}_b(s)\psi(ds) &= \int_0^{\overline{y}} w(y) F_b(dy), \quad \text{and} \\ \int_{\mathcal{S}^n} \widehat{c}_b(s)\psi(ds) &= c_b. \end{aligned} \tag{7}$$

Given contract w, the principal now restricts attention to type spaces in $\mathcal{T}_w^*(F_b, c_b)$, where

$$\mathcal{T}_w^*(F_b, c_b) = \{T \in \mathcal{T}_w^* : \ T \text{ satisfies } (7)\}.$$

In a type space $T \in \mathcal{T}_w^*(F_b, c_b)$, the expected baseline outcome, cost, and contract payments are common knowledge among agents as well as the principal. The baseline action benefits the principal, so we restrict attention to actions (F_b, c_b) such that

$$\int_0^{\overline{y}} y F_b(dy) - c_b \geq 0.$$

Given a baseline action (F_b, c_b), we now show it is still possible for the principal to be weakly worse off from any contract and auction.

Proposition 3. *Fix a baseline action (F_b, c_b). For every contract w, there exists a type space $T \in \mathcal{T}_w^*(F_b, c_b)$ such that $\Pi(T, \sigma \mid a) \leq 0$ for every can-keep auction a and equilibrium $\sigma \in \Sigma(T, a)$.*

Proof. Define $y_b \equiv \int_0^{\overline{y}} y F_b(dy)$ and $w_b \equiv \int_0^{\overline{y}} w(y) F_b(dy)$.

If $w_b - c_b \geq 0$, then we can follow the steps in Lemma 1 and construct a type space $T \in \mathcal{T}_w^*(F_b, c_b)$ such that agents take the baseline action and the principal is weakly worse off in any auction and equilibrium thereof.

If $w_b - c_b < 0$, it does not pay agents to mimic the baseline action. Fix a small enough $y \in [0, \overline{y}]$ so that $y - w(y) \leq y_b - c_b$. Now consider a type space T such that agents produce outcome y at a cost that equals their contract pay $w(y)$. Namely,

$$\widehat{y}_i(s) = y, \quad \widehat{w}_i(s) = w(y), \quad \text{and} \quad \widehat{v}_i(s) = 0$$

for every agent $i \in \{1, \ldots, n\}$ and signal profile $s \in \mathcal{S}^n$. In addition, $\widehat{y}_b(s) = y_b$, $\widehat{c}_b(s) = c_b$ and $\widehat{w}_b(s) = w_b$. T clearly satisfies conditions (1) through (7), so $T \in \mathcal{T}^*_w(F_b, c_b)$. Now consider any can-keep auction a. For any equilibrium $\sigma \in \Sigma(T, a)$, it follows that

$$\begin{aligned}
\Pi(T, \sigma \mid a) &= \int_{\mathcal{S}^n} \int_{\mathbb{R}^n_+} \sum_{i=1}^{n} \Big[q_i(m) \Big(\widehat{y}_i(s) - \widehat{w}_i(s) - (\widehat{y}_b(s) - \widehat{c}_b(s)) \Big) + t_i(m) \Big] \sigma(dm|s) \psi(ds) \\
&= \int_{\mathcal{S}^n} \int_{\mathbb{R}^n_+} \sum_{i=1}^{n} \Big[q_i(m) \Big(y - w(y) - (y_b - c_b) \Big) + t_i(m) \Big] \sigma(dm|s) \psi(ds) \\
&\leq \int_{\mathcal{S}^n} \int_{\mathbb{R}^n_+} \sum_{i=1}^{n} - \Big(q_i(m) \widehat{v}_i(s) - t_i(m) \Big) \sigma(dm|s) \psi(ds) \\
&\leq 0.
\end{aligned}$$

where the last inequality follows because agents may always send message $0 \in \mathcal{M}$ and obtain a non-negative payoff, so they cannot be worse off in equilibrium. ■

Thus, knowledge of the baseline action cannot rule out the possibility that the principal is weakly worse off from any contract and auction. Together with Lemmas 2 and 3, this result implies that a first-price auction of a full-benefit contract is optimal.

But there are many other optimal contracts and auctions here. Indeed, the principal can achieve a payoff guarantee of zero if she sets a reserve price equal to her payoff under the baseline action and shares output with the contracted agent any which way.

Proposition 4. *Fix a baseline action (F_b, c_b) and let $y_b = \int_0^{\overline{y}} y F_b(dy)$. Consider a contract w such that $w(y) \leq y$ for all $y \in [0, \overline{y}]$ and an auction a with a reserve price of $y_b - c_b$. For every type space $T \in \mathcal{T}^*_w(F_b, c_b)$ and equilibrium $\sigma \in \Sigma(T, a)$, $\Pi(T, \sigma \mid a) \geq 0$.*

Proof. Fix a type space $T \in \mathcal{T}^*_w(F_b, c_b)$. T satisfies (1) and $w(y) \leq y$ for all y, so that $\widehat{w}_i(s) \leq \widehat{y}_i(s)$ for all i and $s \in \mathcal{S}^n$. Since auction a has a reserve price of $y_b - c_b$, the principal's revenue is at least $y_b - c_b$ whenever auction a allocates the contract. For any profile of strategies σ, an implication is that

$$\int_{\mathcal{S}^n}\int_{\mathcal{M}^n}\sum_{i=1}^{n}\Big(t_i(m)-q_i(m)(y_b-c_b)\Big)\sigma(dm|s)\psi(ds)\geq 0.$$

In any equilibrium $\sigma \in \Sigma(T,a)$, we conclude that

$$\begin{aligned}\Pi(T,\sigma\,|\,w,a) &= \int\limits_{\mathcal{S}^n}\int\limits_{\mathcal{M}^n}\sum_{i=1}^{n}\Big(q_i(m)\big[\widehat{y}_i(s)-\widehat{w}_i(s)-(y_b-c_b)\big]+t_i(m)\Big)\sigma(dm|s)\psi(ds)\\ &\geq \int\limits_{\mathcal{S}^n}\int\limits_{\mathcal{M}^n}\sum_{i=1}^{n}q_i(m)\big[\widehat{y}_i(s)-\widehat{w}_i(s)\big]\sigma(dm|s)\psi(ds)\\ &\geq 0.\end{aligned}$$

■

3.3 Other Known Actions

In section 2, the principal knows agents may take the baseline action but ignores what other actions are available. We now allow the principal to know of other available actions.

As in Carroll (2015), the principal knows that agents may choose among a compact set of actions $\mathcal{A}\subseteq\Delta([0,\overline{y}])\times[0,\infty)$ and ignores what other actions are also available. A type space T is consistent with a set of actions $\mathcal{A}$ under contract w if

$$\widehat{v}_i(s)\geq\max_{(F,c)\in\mathcal{A}}\int_0^{\overline{y}}w(y)F(dy)-c\quad\text{for all } i\in\{1,\ldots,n\}\text{ and } s\in\mathcal{S}^n. \tag{8}$$

Given contract w, the principal considers type spaces in $\mathcal{T}_w^*(\mathcal{A})$, where

$$\mathcal{T}_w^*(\mathcal{A})=\{T\in\mathcal{T}_w^*:\ T\text{ satisfies (8)}\}.$$

Note that we have not included the baseline action in set $\mathcal{A}$: a type space $T\in\mathcal{T}_w^*(\mathcal{A})$ need not be consistent with any given baseline action (F_b,c_b). Our principal therefore continues to ignore its consequences.

Here, the principal continues to allow for the possibility that agents mimic the baseline action — it might be better than any action in $\mathcal{A}$. As a result (Lemma 1), the principal's payoff guarantee is still non-positive. Moreover, a first-price auction of a full-benefit contract is optimal and other contracts are not, regardless of the auction pairing. But I don't think it makes much sense to think that the principal knows about other actions available to agents yet ignores the baseline action. A more interesting extension would have the principal know *the baseline action as well as other actions available to agents. This is the extension I am struggling with.*

4 Discussion

We have presented a class of mechanisms — auctions of contracts — as an alternative for a principal to hire an agent when she faces an opportunity cost of contracting and is unsure about its magnitude. The principal's opportunity cost is given by her payoff under the baseline action, the action she would take if she did not hire anyone. When the baseline action is also available to agents, we showed that the principal obtains the best possible guarantee by selling a contract that pays all her benefit from the realized outcome with a first-price auction.

Our results speak to existing work on robustness in contract design. We have shown that contracts other than the full-benefit contract can make our principal worse off, regardless of the auction pairing. This finding contrasts with the well-known positive payoff guarantees of linear contracts (Chassang 2013, Carroll 2015), which depend on the assumption that the principal's baseline action produces no output. This assumption may make sense or not, depending on the application of interest. Our marginal contribution to this literature is to show how the principal can still hire while avoiding negative payoffs when this assumption is not warranted.

At a broader level, we contribute to the literature on the design of mechanisms that perform well in a wide range of situations. The joint contract and auction design problem forces us to simultaneously address designer uncertainty about the actions and the information available to agents. Our approach defines type spaces to include both a common prior information structure that specifies a hierarchy of agents' beliefs about their values of the contract, as well as a specification of the counterfactual outcomes that agents would produce under contract and that the principal would produce if she did not hire anyone. The principal considers a class of contract-specific type spaces and judges contract and auction pairs according to her payoff in the worst possible type space. This approach provides the interface we need to draw insights from the fertile literatures on informationally robust auction design (Bergemann, Brooks, and Morris 2017, Bergemann, Brooks, and Morris 2019, Bergemann, Brooks, and Morris 2020, Brooks and Du 2021) and robust contract design (Carroll 2015).

That said, the optimal payoff guarantee for our principal is zero. So, why bother auctioning off a contract? Our motivation for this paper is not to guide a principal who literally values a contract and an auction according to their worst possible performance — she would certainly be indifferent between selling a contract and taking the baseline action, at best. Rather, we have aimed to find contracts and auctions that effectively outsource production under few assumptions. This is valuable for a principal who somehow needs the approval of others: the assumptions we lay out are a consensus that people must reach to cast the principal's decisions in a favorable light.

References

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2017). "First-Price Auctions With General Information Structures: Implications for Bidding and Revenue". In: *Econometrica* 85.1, pp. 107–143. DOI: https://doi.org/10.3982/ECTA13958. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13958. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA13958.

— (May 2019). "Revenue Guarantee Equivalence". In: *American Economic Review* 109.5, pp. 1911–29. DOI: 10.1257/aer.20180773. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20180773.

— (2020). "Countering the winner's curse: Optimal auction design in a common value model". In: *Theoretical Economics* 15.4, pp. 1399–1434. DOI: https://doi.org/10.3982/TE3797. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/TE3797. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/TE3797.

Brooks, Benjamin and Songzi Du (2021). "Optimal Auction Design With Common Values: An Informationally Robust Approach". In: *Econometrica* 89.3, pp. 1313–1360. DOI: https://doi.org/10.3982/ECTA16297. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA16297. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16297.

Bulow, Jeremy and Paul Klemperer (2002). "Prices and the Winner's Curse". In: *The RAND Journal of Economics* 33.1, pp. 1–21. ISSN: 07416261. URL: http://www.jstor.org/stable/2696372 (visited on 05/17/2024).

Carroll, Gabriel (Feb. 2015). "Robustness and Linear Contracts". In: *American Economic Review* 105.2, pp. 536–63. DOI: 10.1257/aer.20131159. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20131159.

— (2019). "Robustness in Mechanism Design and Contracting". In: *Annual Review of Economics* 11.1, pp. 139–166. DOI: 10.1146/annurev-economics-080218-025616. eprint: https://doi.org/10.1146/annurev-economics-080218-025616. URL: https://doi.org/10.1146/annurev-economics-080218-025616.

Chassang, Sylvain (2013). "Calibrated Incentive Contracts". In: *Econometrica* 81.5, pp. 1935–1971. DOI: https://doi.org/10.3982/ECTA9987. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA9987. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9987.

Engelbrecht-Wiggans, Richard, Paul R. Milgrom, and Robert J. Weber (1983). "Competitive bidding and proprietary information". In: *Journal of Mathematical Economics* 11.2, pp. 161–169. ISSN: 0304-4068. DOI: https://doi.org/10.1016/0304-4068(83)90034-4. URL: https://www.sciencedirect.com/science/article/pii/0304406883900344.

Kremer, Michael and Rachel Glennerster (2004). *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton University Press. URL: http://www.jstor.org/stable/j.ctt1dr365r.

Kremer, Michael, Jonathan Levin, and Christopher M. Snyder (May 2020). "Advance Market Commitments: Insights from Theory and Experience". In: *AEA Papers and Proceedings* 110, pp. 269–73. DOI: 10.1257/pandp.20201017. URL: https://www.aeaweb.org/articles?id=10.1257/pandp.20201017.

— (2022). "Designing Advance Market Commitments for New Vaccines". In: *Management Science* 68.7, pp. 4786–4814. DOI: 10.1287/mnsc.2021.4163. eprint: https://doi.org/10.1287/mnsc.2021.4163. URL: https://doi.org/10.1287/mnsc.2021.4163.