2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Information Asymmetry and Contract Design with Applications to Agriculture

Alessandro Bonatti Munther Dahleh Thibaut Horel Mardavij Roozbehani

***Abstract*— We consider situations in which an intermediary facilitates the interactions of one or several players with a downstream market in a game of incomplete information. Our key assumption is the presence of information asymmetries: while the intermediary has in general better (less noisy) information about the observable parameters of the game, the players have better information about their own private parameters and preferences. The intermediary seeks to influence the actions of the agents by offering side information and/or certain guarantees regarding the outcomes. For instance, farmers aim to sell their produce in a downstream market. The intermediary has access to more accurate signals regarding the downstream market (e.g. demand, prices, etc.), while the farmers are aware of their own private cost structure. The central problem is then to understand how to design contracts for exchanging information and mediating the interaction between the players and the downstream market in such a way that generates value to the players and revenue to the intermediary.**

Prior work on information design with elicitation has shown that in the presence of competition between the players, the intermediary can generate value by coordinating the players' actions in such a way that reduces the negative externalities they exert on each other. In this work, we focus on the question of risk-aversion. Our first result is negative and shows that the intermediary cannot generate revenue when interacting with a single risk-neutral or risk-seeking player. We then explore how this result changes under various relaxations of the model which include altering the risk preference of the player or changing the timing of the game.

I. Introduction

Information design is the study of what information to reveal, when to reveal it, and to whom in order to achieve certain goals [1]. Applications abound in a variety of domains including smart infrastructure systems, game theory and economics, finance, autonomy and computer science. Interest in information design has grown in the control and systems community in recent years due to its inherent ties with decision theory. Recent papers have studied information design in the context of scheduling games [2], routing games [3], congestion control in transportation networks [4], [5], and coordination of autonomous systems [6].

Related to our work is also the field of *contract design*, which studies how to design contracts between participating agents to achieve desirable outcomes [7]. Contract design, like information design, has also garnered significant attention from decision and control scientists in recent years. It has found notable applications in areas such as energy regulation and procurement [8], [9], [10], [11], as well as in cyber-physical security [12].

Motivated by the much less studied application to farming and agricultural markets, we adopt in this paper a more abstract view of problems that arise due to the interplay between information design and contract design. We focus on a simple model involving a single agent choosing an action (e.g. a farmer choosing which crop to plant) that is directly or indirectly valued in a downstream market. The exact properties of this market (e.g. demand, prices, etc.) are unknown to the agent but are observed or accurately forecast by the intermediary, henceforth called the *principal* in accordance with the terminology in contract theory. The main problem that the principal faces in such a situation is to design a contract specifying what information to exchange with the agent, at which price, and how to compensate the agent based on the outcome and/or their action. Our main contribution is to formalize and compare two classes of contracts that can be used to solve the principal's problem:

- In the first class, the principal acts as an intermediary, buying the agent's production at a guaranteed price and selling it in the downstream market. In this sense, the intermediary is acting as an insurer. We show that there exists revenue-generating contracts within this class if and only if the agent is risk averse.
- In the second class, the principal acts as an information seller, collecting payments in exchange for an informative signal about the downstream market, but does not interact with the downstream market themselves. In contrast to the first class, we show that revenue-generating contracts always exist in this class.

Related work

As already mentioned, our work draws from two lines of work: contract theory [7] and information design [1]. A significant portion of the literature in contract theory is concerned with an information asymmetry different from the one we study here, due to the inability for the principal to observe the agent's action and preferences, resulting in a problem known as *moral hazard* (standard textbooks on this topic include [13], [14]). The case of an informed principal (who observes a signal before offering a contract) has been comparatively less studied with some of the seminal exceptions including [15], [16]. Information design when a signal is sent to a single player is also known as *Bayesian persuasion* [17], [18], although the setting has also been extended to situations with multiple signal senders [19]. The interaction between moral hazard—which we do not study in the present work—and Bayesian persuasion has been studied in [20]. The work in Bayesian persuasion differs from the present work in that it usually does not involve payments: the focus is on studying how to influence the actions of one or multiple agents solely via informative signals. In contrast, our contracts involve a payment that can depend either on the action taken by the agent—when the principal acts as an intermediary—or on the information being sold—when the principal acts as an information seller.

II. Model

A. Downstream market

The parameters of the downstream market are represented by a state ω drawn from a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The agent's action space is $\mathcal{A}$, and for an action $a \in \mathcal{A}$ and state $\omega \in \Omega$, the market value for the agent's action is written as $v(a;\omega)$, where $v : \mathcal{A} \times \Omega \to \mathbb{R}$.

The agent's cost is described by a function $c : \mathcal{A} \to \mathbb{R}$ and we assume that their utility is an increasing function of value less cost. Specifically, given action a and state ω, the agent's utility is

$$\tilde{u}(a;\omega) := u\big(v(a;\omega) - c(a)\big),$$

where $u : \mathbb{R} \to \mathbb{R}$ is an increasing function. When u is concave the agent is risk-averse, whereas the agent is risk-seeking when u is convex. Because u is increasing, maximizing the agent's utility for a fixed realization $\omega \in \Omega$ is equivalent to maximizing $v(a;\omega) - c(a)$. The following assumption guarantees that this optimization problem is well-posed.

Assumption 1. For every $\omega \in \Omega$, the function $a \mapsto v(a;\omega) - c(a)$ attains a maximum over $\mathcal{A}$. We denote by a_ω an action at which the maximum is reached when the state is ω, with ties broken arbitrarily when the maximum is not unique.

Example 2. Consider an agricultural market in which K produces can be sold at marginal prices $\omega := (p_1, \ldots, p_k) \in \mathbb{R}^K_{\geq 0}$. The agent is a farmer that produces quantities $a := (q_1, \ldots, q_k)$ subject to a capacity constraint Q. That is, the action space is

$$\mathcal{A} := \Big\{(q_1, \ldots, q_K) \in \mathbb{R}^K_{\geq 0} \,\Big|\, \sum_{k=1}^{K} q_k \leq Q\Big\}$$

We assume that the farmer's production is small enough not to affect prices and that the farmer's marginal production cost for product k is constant and equal to c_k. Hence, the market value and cost functions for action $a \in \mathcal{A}$ and state ω are given by

$$v(a;\omega) = \sum_{k=1}^{K} p_k q_k \quad \text{and} \quad c(a) = \sum_{k=1}^{K} c_k q_k.$$

Since $v(a;\omega) - c(a)$ is linear in a, we can always find a maximum of the farmer's utility at an extreme point of $\mathcal{A}$. Specifically, let k be the index of a product that maximizes the marginal profit $p_k - c_k$. If $p_k - c_k$ is non-negative then an optimal action is $a_\omega = e_k$, where e_k is the kth standard basis vector. Otherwise, the optimal action is to not produce anything ($a_\omega = 0$).

Example 3. In this example, there is unique produce and the farmer's action is interpreted as the choice of either how much to produce or the quality at which to produce. In both cases, we have $\mathcal{A} = \mathbb{R}_{\geq 0}$ and we model the dependency of the market's marginal value on quantity/quality as linear,

$$v(a;\omega) = a(\lambda + \mu \cdot a).$$

where the state $\omega := (\lambda, \mu)$ are the parameters of this linear dependency. In the interpretation where a is quantity, we have $\mu < 0$, whereas $\mu > 0$ when a is quality.

The farmer's cost is quadratic, $c(a) = -a^2$. Assuming that $\mu < 1$, $v(a;\omega) - c(a)$ is concave in a and the farmer's utility reaches a unique maximum at

$$a_\omega = \frac{\lambda}{2(1-\mu)}.$$

B. Information structure

The distribution $\mathbb{P}$ is known by both players and acts as a common prior about the state ω. We assume that the principal observes ω before any interaction with the agent takes place. This could serve as a simplified model for the situation in which the principal can forecast the outcome relatively accurately, with significantly less error than the agent can. The agent is a Bayesian utility maximizer: whenever they observe a signal S, they update their belief about ω and compute their utility as the conditional expectation $\mathbb{E}[\tilde{u}(a;\omega) \mid S]$.

The principal does not take any action and only transacts with the agent and the downstream market. The principal is risk-neutral[1], hence their utility is simply the sum of the monetary transfers with the agent and the downstream market. The exact form it takes depends on the specifics of the contract negotiated with the agent and will be explicated in Section III.

III. Contracts

A. Principal as an "insurer"

Here, we explore the case where the principal acts as an insurer in the following sense. Upon observing (or obtaining an accurate estimate of) ω, the principal declares $\tilde{\omega}$ (random variable correlated with ω) to the agent and offers $v(a;\tilde{\omega})$ for any action a. In effect, this creates a "de-risked" market in which the agent has guaranteed utility $\tilde{u}(a;\tilde{\omega}) := u\big(v(a;\tilde{\omega}) - c(a)\big)$ for taking action a, compared to expected utility $\mathbb{E}[\tilde{u}(a;\omega)]$ in the actual downstream market.

Consequently, the agent accepts the offer if the following participation constraint holds

$$\max_{a\in\mathcal{A}} \tilde{u}(a;\tilde{\omega}) \geq \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid \tilde{\omega}] \tag{1}$$

If the agent accepts the offer, the agent takes action $a_{\tilde{\omega}}$ that maximizes the left-hand side and is guaranteed to exist by Assumption 1. The principal gets to sell the production of the agent for $v(a_{\tilde{\omega}};\omega)$ resulting in profit

$$v(a_{\tilde{\omega}};\omega) - v(a_{\tilde{\omega}};\tilde{\omega})$$

If the agent rejects the offer, they play directly in the downstream market with action a that maximizes $\mathbb{E}[\tilde{u}(a;\omega) \mid \tilde{\omega})]$, and the principal makes no profit.

[1]All of our results and analyses extend to the case where the principal is risk seeking.

Remark. The crucial point is that the offer $\tilde{\omega}$ depends on ω and thus gives a signal to the agent about the underlying state of the world. Consequently, when the agent is assessing their outside option, their belief about ω is the posterior distribution conditional on $\tilde{\omega}$.

Proposition 4. *If the agent is risk-neutral or risk-seeking, i.e., if $u(\cdot)$ is convex, the principal's expected profit is non-positive, regardless of the strategy used.*

Proof. If the participation constraint does not hold, then the principal makes no profit. So we focus on the case where it holds. By (1) and using the definition of $a_{\tilde{\omega}}$:

$$\begin{aligned}\tilde{u}(a_{\tilde{\omega}};\tilde{\omega}) &= \max_{a\in\mathcal{A}} \tilde{u}(a;\tilde{\omega}) \\ &\geq \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)\mid\tilde{\omega}] \geq \mathbb{E}[\tilde{u}(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}].\end{aligned}$$

We lower bound the right-hand side by using Jensen's inequality for conditional expectations and obtain

$$u\big(v(a_{\tilde{\omega}};\tilde{\omega}) - c(a_{\tilde{\omega}})\big) \geq u\left(\mathbb{E}[v(a_{\tilde{\omega}};\omega) - c(a_{\tilde{\omega}})\mid\tilde{\omega}]\right).$$

By monotonicity of $u(\cdot)$ and linearity of conditional expectation, the above is equivalent to:

$$v(a_{\tilde{\omega}};\tilde{\omega}) - c(a_{\tilde{\omega}}) \geq \mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}] - \mathbb{E}[c(a_{\tilde{\omega}})\mid\tilde{\omega}],$$

or alternatively:

$$\mathbb{E}[c(a_{\tilde{\omega}})\mid\tilde{\omega}] - c(a_{\tilde{\omega}}) \geq \mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}] - v(a_{\tilde{\omega}};\tilde{\omega}).$$

Finally, by taking expectations on both sides and invoking the law of iterated expectations we have:

$$0 \geq \mathbb{E}\big[v(a_{\tilde{\omega}};\omega)\big] - \mathbb{E}\big[v(a_{\tilde{\omega}};\tilde{\omega})\big],$$

where, we recognize the principal's expected profit on the right-hand side. □

Proposition 5. *If the agent is strictly risk-averse, i.e., if $u(\cdot)$ is strictly concave (and increasing), it is possible for the principal to make positive profit in expectation.*

Proof. To simplify the notation and for clarity, let us consider the case where the cost function $c(\cdot)$ is identically zero. The proof for the general case is similar. The participation constraint is

$$u\big(v(a_{\tilde{\omega}};\tilde{\omega})\big) \geq \mathbb{E}\big[u\big(v(a_{\tilde{\omega}};\omega)\big)\;\big|\;\tilde{\omega}\big] \tag{2}$$

whereas, Jensen's inequality works in the other direction (comparing with proof of proposition 4) due to concavity of $u(\cdot)$. Therefore,

$$\delta := u\big(\mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}]\big) - \mathbb{E}\big[u\big(v(a_{\tilde{\omega}};\omega)\big)\;\big|\;\tilde{\omega}\big] \geq 0.$$

When ω is not uniquely determined from $\tilde{\omega}$, i.e., when $\tilde{\omega}$ is not a deterministic invertible function of ω, the above inequality is strict due to strict concavity of $u(\cdot)$, and $\delta > 0$ provides a positive gap that can be exploited. It then follows that the quantity

$$\delta' = \delta'(\tilde{\omega}) := \mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}] - u^{-1}\big(\mathbb{E}\big[u\big(v(a_{\tilde{\omega}};\omega)\big)\;\big|\;\tilde{\omega}\big]\big) \tag{3}$$

is also strictly positive. The principal can then offer the following payment structure (contract) to the agent:

$$v(a_{\tilde{\omega}};\tilde{\omega}) = \mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}] - \eta, \quad \eta\in(0,\delta') \tag{4}$$

$$> u^{-1}\big(\mathbb{E}\big[u\big(v(a_{\tilde{\omega}};\omega)\big)\;\big|\;\tilde{\omega}\big]\big) \tag{5}$$

where the inequality follows from the fact that δ' is strictly positive and $\eta\in(0,\delta')$. Therefore, due to monotonicity of $u(\cdot)$, the participation constraint (1) is satisfied and the agent always participates. On the other hand, the principal's profit is:

$$v(a_{\tilde{\omega}};\omega) - v(a_{\tilde{\omega}};\tilde{\omega}) \tag{6}$$

which, in expectation is equal to $\eta > 0$. □

Remark. Note that in this case the principal is able to make profit by assuming a risk that the agent is not willing to take. Therefore, the principal's profit can be positive only in expectation, and not almost surely. This intuition is confirmed by the proof of Proposition 5. If (6) is positive almost surely, then the agent's participation constraint (2) cannot be satisfied. Furthermore, the result relies on strictness of Jensen's inequality. Therefore, in addition to strict concavity of the agent's utility, we need uncertainty in the conditional universe where $\tilde{\omega}$ is given. If the contract and/or the signal $\tilde{\omega}$ fully reveal ω, the situation is effectively fully de-risked for the agent, and the principal cannot make profit per Proposition 4. Lastly, we remark that the contract value $v(a_{\tilde{\omega}};\tilde{\omega})$ cannot be an unbiased estimator of $\mathbb{E}[v(a_{\tilde{\omega}};\omega)\mid\tilde{\omega}]$, otherwise, by the law of total expectations the principal cannot make profit. Thus, presence of a bias term like η as in Equation (4) is necessary. Although the bias term in (4) is not unique and can take other forms.

A lower bound on the principal's profit: Consider now the special case where the principal's signal is non-informative, e.g., is always constant or independent of ω. Thus, the agent's optimal action is independent of the signal $\tilde{\omega}$, and the only decision the agent makes is whether to take the outside option (sell in the downstream market) or accept the principal's contract. Since $\tilde{\omega}$ is non-informative, the agent either always participates in the contract or never does. In order to make positive profit then the principal must ensure that the participation constraint is always satisfied. Let us use $\tilde{\omega} = \emptyset$ to denote the non-informative signal and let $\bar{a} = a_{\tilde{\omega}} = a_{\emptyset}$ denote the (constant) action of the agent in this case. The principal's expected revenue from the market is then $\mathbb{E}[v(\bar{a};\omega)]$, and depends only on the statistics of the market or state of the word $\omega\in\Omega$. In order to maximize profit the principal must minimize the expected payment to the agent while guaranteeing participation. Equivalently, the principal solves the following optimization problem:

$$\begin{aligned}\inf \quad & v(\bar{a};\emptyset) \\ \text{s.t.} \quad & u(v(\bar{a};\emptyset)) > \mathbb{E}[u(v(\bar{a};\omega))]\end{aligned} \tag{7}$$

It can be verified that the infimal value of the above optimization problem is

$$\gamma = u^{-1}\big(\mathbb{E}[u(v(\bar{a};\omega))]\big)$$

The principal's expected profit is thus lower bounded by $\mathbb{E}[v(\bar{a};\omega)] - u^{-1}\big(\mathbb{E}[u(v(\bar{a};\omega))]\big)$, which is non-negative by Jensen's inequality. Note that this lower bound is equal to $\delta'(\emptyset)$, where $\delta'(\tilde{\omega})$ is defined in (3).

B. Principal as an information seller

In this section, we turn to a different class of contracts in which the principal acts only as an information seller not as an insurer. Specifically, a contract in this case is a triplet $(\mathcal{S}, \sigma, t)$ where:

- $\mathcal{S}$ is the signal space.
- $\sigma : \Omega \to \Delta(\mathcal{S})$ is the signalling function mapping a state ω to a distribution $\sigma(\omega)$ over the signal space.
- t is the payment made by the agent to the principal.

The interpretation of a such a contract in natural language is: *in exchange for a payment of t, the principal sends to the agent a signal $S \in \mathcal{S}$ distributed according to $\sigma(\omega)$.*

Participation constraint: Let us now examine more closely the condition under which such a contract is profitable for the agent.

If the contract is accepted, upon receiving the signal S, the agent updates their belief about the unknown state ω and chooses an action that maximizes the conditional expectation $\mathbb{E}[\tilde{u}(a;\omega) \mid S]$. Taking an expectation over all possible realizations of the signal S yields the expected utility that the agent believes is achievable by opting in:

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right].$$

If the contract is rejected, the agent does not receive any information from the principal and simply chooses an action that maximizes their expected utility under the prior distribution of ω, resulting in utility

$$\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)].$$

Consequently, the agent (weakly) prefers to accept the contract if the following inequality holds

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] - t \geq \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)]. \tag{8}$$

An important difference with the participation constraint (1) is that there is no conditioning in the expectation on the right-hand side of (8). Indeed, when deciding whether to opt in the contract or not, the agent has no information to condition on and can only reason about the perceived *informativeness* of the signalling function σ. Here, informativeness should not be understood in an information-theoretic sense but rather in the utilitarian sense of allowing the agent to make better decisions in the downstream market. Informativeness in this sense is precisely quantified by the gap

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] - \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)],$$

and the participation constraint can equivalently be stated as requiring the payment t to be no more than this gap.

Proposition 6. *For each signalling function σ, there exists a non-negative payment t that satisfies the participation constraint* (8)*.*

A contract that maximizes the principal's revenue is to reveal the true state ω in exchange for the payment

$$t_{\max} := \mathbb{E}\Big[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega)\Big] - \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)].$$

Proof. Let $\sigma : \Omega \to \Delta(\mathcal{S})$ be a signalling function and let S be a random variable distributed according to $\sigma(\omega)$ where ω is drawn from the prior $\mathbb{P}$. Then for each $a \in \mathcal{A}$ we have by definition of the supremum

$$\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S] \geq \mathbb{E}[\tilde{u}(a;\omega) \mid S]$$

almost surely over the realization of S. Taking expectations on both sides yields

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] \geq \mathbb{E}\big[\mathbb{E}[\tilde{u}(a;\omega \mid S]\big] = \mathbb{E}[\tilde{u}(a;\omega)],$$

where the equality uses the law of total expectation. Finally, taking a supremum over $a \in \mathcal{A}$ on the right-hand side gives

$$\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] \geq \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)].$$

Define the payment t_σ by

$$t_\sigma := \mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] - \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)].$$

Then, the previous inequality immediately implies that $t_\sigma \geq 0$ and that the participation constraint (8) holds for t_σ. This concludes the proof of the first claim.

For the second claim, observe that for the payment t_σ defined above, the participation constraint is in fact binding. In other words, t_σ is the maximum payment that can be collected when using the signalling function σ while still guaranteeing participation. Consequently, a revenue-maximizing contract is obtained by maximizing t_σ over all possible signalling functions σ. For each $a \in \mathcal{A}$ and $\omega \in \Omega$, we have by definition of the maximum

$$\tilde{u}(a;\omega) \leq \max_{a\in\mathcal{A}} \tilde{u}(a;\omega),$$

where the maximum is reached by Assumption 1. By the positivity of conditional expectations, this implies

$$\mathbb{E}[\tilde{u}(a;\omega) \mid S] \leq \mathbb{E}\Big[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega) \;\Big|\; S\Big].$$

Taking a supremum over $a \in \mathcal{A}$ on the left-hand side yields

$$\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S] \leq \mathbb{E}\Big[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega) \;\Big|\; S\Big].$$

Finally, taking expectations on both sides gives

$$\begin{aligned}\mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] &\leq \mathbb{E}\left[\mathbb{E}\Big[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega) \;\Big|\; S\Big]\right] \\ &= \mathbb{E}\Big[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega)\Big],\end{aligned}$$

where the equality follows from the law of total expectation. Consequently we just established that for any signalling function σ,

$$\begin{aligned} t_\sigma &:= \mathbb{E}\left[\max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega) \mid S]\right] - \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)] \\ &\leq \mathbb{E}\left[\max_{a\in\mathcal{A}} \tilde{u}(a;\omega)\right] - \max_{a\in\mathcal{A}} \mathbb{E}[\tilde{u}(a;\omega)]. \end{aligned}$$

Furthermore, it is easy to see that equality is reached with $S = \omega$ for which $\mathbb{E}[\tilde{u}(a;\omega) \mid S] = \tilde{u}(a;\omega)$. This concludes the proof of the second claim. □

IV. Discussion

The two classes of mechanisms studied respectively in Section III-A and Section III-B reveal two distinct ways in which information can be used to provide value to an uninformed agent and in turn generate revenue:

1) In the first case, information is used as a way to reduce the agent's uncertainty about the unknown state. The principal presents themselves to the agent as an intermediate and less risky market that lies in between the original market and the agent. In such a situation, one could reasonably expect that the agent's attitude toward risk is the primary factor shaping their interaction with the principal. This is precisely what we found in Proposition 4 and Proposition 5, showing that risk aversion is a necessary and sufficient condition for the existence of revenue-generating contracts.
2) In the second case, the value of information is derived from the improvement in decision-making abilities that it induces. This places the principal in the position of a seller that designs an information good in the form of a signal. Importantly, the signal must be acquired by the agent in exchange for a payment *before* it can be observed. In such a situation, we found in Proposition 6 that it is always possible for the principal to generate revenue, with the optimum revenue achieved by revealing all the information to the agent. This can be understood as another manifestation of the "information never hurts" principle.

There are of course many directions left to explore for future work. First of all, it would be interesting to quantify more precisely the difference between the two aforementioned cases: when is it more profitable or more valuable to the agent to have the principal act as a "insurer" rather than as an information seller? Next, our model could be extended to situations in which the agent's cost is of the form $c(a;\theta)$, where the parameter θ is the agent's privately observed type. The contracts would now need to take the form of a *mechanism*—as studied in the economics fields of auction theory and mechanism design—that would truthfully elicit the agent's type or let them self-select into a "menu" of options. Finally, all the above questions could also be asked in environment with multiple strategic agents. In such environments, information can be used to coordinate the agents' actions—which has implications for how to sell it as recently studied in [21]—and allows for the design of risk-pooling contracts.

References

[1] D. Bergemann and S. Morris, "Information design: A unified perspective," *Journal of Economic Literature*, vol. 57, no. 1, pp. 44–95, 2019.
[2] N. Heydaribeni and K. Savla, "Information design for a non-atomic service scheduling game," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 3800–3505.
[3] Y. Zhu and K. Savla, "Information design in nonatomic routing games with partial participation: Computation and properties," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 2, pp. 613–624, 2022.
[4] S. Das, E. Kamenica, and R. Mirka, "Reducing congestion through information design," in *2017 55th annual allerton conference on communication, control, and computing (allerton)*. IEEE, 2017, pp. 1279–1284.
[5] M. Wu and S. Amin, "Information design for regulating traffic flows under uncertain network state," in *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2019, pp. 671–678.
[6] M. Le Treust and T. Tomala, "Information design for strategic coordination of autonomous devices with non-aligned utilities," in *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2016, pp. 233–242.
[7] P. Bolton and M. Dewatripont, *Contract theory*. MIT press, 2004.
[8] D. G. Dobakhshari and V. Gupta, "Optimal contract design for incentive-based demand response," in *2016 American Control Conference (ACC)*, 2016, pp. 3219–3224.
[9] ——, "A contract design approach for phantom demand response," *IEEE Transactions on Automatic Control*, vol. 64, no. 5, pp. 1974–1988, 2019.
[10] M. Balandat, F. Oldewurtel, M. Chen, and C. Tomlin, "Contract design for frequency regulation by aggregations of commercial buildings," in *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2014, pp. 38–45.
[11] H. Tavafoghi and D. Teneketzis, "Optimal contract design for energy procurement," in *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2014, pp. 62–69.
[12] J. Chen and Q. Zhu, "A linear quadratic differential game approach to dynamic contract design for systemic cyber risk management under asymmetric information," in *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2018, pp. 575–582.
[13] J.-J. Laffont and D. Martimort, *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, 2002. [Online]. Available: http://www.jstor.org/stable/j.ctv7h0rwr
[14] P. Bolton and M. Dewatripont, *Contract Theory*. The MIT Press, December 2004.
[15] E. Maskin and J. Tirole, "The principal-agent relationship with an informed principal, II: Common values," *Econometrica*, vol. 60, no. 1, pp. 1–42, 1992. [Online]. Available: http://www.jstor.org/stable/2951674
[16] R. B. Myerson, "Mechanism design by an informed principal," *Econometrica*, vol. 51, no. 6, pp. 1767–1797, 1983. [Online]. Available: http://www.jstor.org/stable/1912116
[17] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, October 2011. [Online]. Available: https://www.aeaweb.org/articles?id=10.1257/aer.101.6.2590
[18] E. Kamenica, "Bayesian persuasion and information design," *Annual Review of Economics*, vol. 11, pp. 249–272, 2019.
[19] M. Gentzkow and E. Kamenica, "Competition in persuasion," *The Review of Economic Studies*, vol. 84, no. 1, p. 300–322, 10 2016.
[20] R. Boleslavsky and K. Kim, "Bayesian persuasion and moral hazard," *Available at SSRN 2913669*, 2018.
[21] A. Bonatti, M. Dahleh, T. Horel, and A. Nouripour, "Coordination via selling information," 2023.