Selling Information in Competitive Environments*

Alessandro Bonatti† Munther Dahleh‡ Thibaut Horel† Amir Nouripour†

Abstract

Data buyers compete in a game of incomplete information about which a single data seller owns some payoff-relevant information. The seller faces a joint information- and mechanism-design problem: deciding which information to sell, while eliciting the buyers' types and imposing payments. We derive the welfare- and revenue-optimal mechanisms for a class of games with binary actions and states. Our results highlight the critical properties of selling information in competitive environments: (i) the negative externalities arising from buyer competition increase the profitability of recommending the correct action to one buyer exclusively; (ii) for the buyers to follow the seller's recommendations, the degree of exclusivity must be limited; (iii) the buyers' obedience constraints also reduce the distortions in the allocation of information introduced by a monopolist; (iv) as competition becomes fiercer, these limitations become more severe, weakening the impact of market power on the allocation of information.

Keywords: Data, Competition, Screening, Information Design, Externalities.

JEL Codes: D43, D82, D83.

*Bonatti acknowledges financial support through NSF Grant SES–1948692. Dahleh, Horel, and Nouripour acknowledge financial support from OCP Africa. We thank the Associate Editor, two anonymous referees, Stephen Morris, Maryann Rui, and seminar participants at UChicago, UIUC, INFORMS, Harvard, and MIT for insightful comments and valuable feedback that helped greatly improve the paper.

†Massachusetts Institute of Technology, Sloan School of Management

‡Massachusetts Institute of Technology, Institute for Data, Systems, and Society

1 Introduction

Markets for information shape a growing fraction of the economy. Information is sold directly (e.g., credit bureaus sell consumer scores to lenders, and research institutions sell data to financial traders) but also indirectly (e.g., digital platforms offer advertisers access to a targeted audience, and hedge funds sell shares of the portfolios they build based on superior information).[1] The allocation of information affects the distribution of market power in the downstream markets where that information is used, thereby critically impacting consumer and social surplus. Understanding how these markets work, and what is special about them, is then a first-order economic and social issue.

In this paper, we study how private information and buyer competition interact in determining the optimal allocation and price of information. Our objective is threefold: (i) to provide qualitative insights into the structure of the revenue-maximizing mechanisms for the sale of information, (ii) to determine how information differs from physical goods in this respect, and (iii) to assess the impact of market power in the information sector on competition in downstream markets. We propose a tractable formulation of this problem where the profitability of acquiring information for any buyer is unknown to the seller (e.g., buyers have private cost, asset holdings, risk preferences), and buyers of information compete with one another (e.g., financial traders, advertisers, lenders). We cast the monopolist seller's choice of a mechanism for the sale of information as an *information design* problem *with elicitation* (Bergemann and Morris, 2019). We consider finitely many data buyers and a data seller. The data buyers compete in a simultaneous-moves, finite game of incomplete information (the "downstream game"). The monopolist is informed about a payoff-relevant state variable and sells informative signals to the buyers. Each buyer also has a payoff type in the downstream game, i.e., their willingness to pay for any signal is their private information. Thus, the seller must first elicit the buyers' private payoff types and then sell them informative signals. As these signals can be viewed as action recommendations, the seller faces a joint mechanism and information design problem, wherein their choice of information structure is subject to the buyers' obedience and truthful reporting constraints.

More specifically, a direct mechanism maps the state of the world and the buyers' reported types into a distribution over informative signals and payments. An important property of our setting is that the seller can design any statistical experiment but lacks complete control over the buyers' actions. This is because information is only valuable insofar as it affects behavior (Blackwell, 1953), and the buyers retain control over their downstream actions.

[1] We describe the targeted advertising example at length below. See Admati and Pfleiderer (1990) and Bergemann and Bonatti (2019) for a discussion of direct vs. indirect sales of information.

Likewise, the seller can design the information revealed to any buyer when one or more buyers do not participate in the mechanism. However, the seller cannot prevent any buyer from playing in the downstream game under their prior information only. Therefore, the designer has partial but not full control over each buyer's outside options, which partially relaxes the buyers' participation constraints.

Main Results We begin with a characterization of the seller's constraints in Section 3. This characterization holds whenever the buyers' payoffs are linear in their private type. For such payoff structures, we show that incentive compatibility of the mechanism is equivalent to *separately* incentivizing truthful reporting and obedience of the buyers. In other words, double deviations—wherein a buyer both misreports their type and deviates from the seller's action recommendation—are no more profitable to the buyers than one-shot deviations.

Next, we focus on the simplest instance of this complex problem—a binary-action downstream game of incomplete information where the state identifies which action is dominant for each data buyer. In this game, acquiring information imposes negative externalities on the other buyers: the better informed a buyer is about the state, the *lower* the resulting payoff for the other buyers. In other words, the seller designs a mechanism in the presence of externalities stemming from the competition among buyers (Jehiel and Moldovanu, 2006).

We then turn to the welfare-optimal mechanism for the allocation of information and the revenue-maximizing mechanism for a monopolist seller. Our results highlight two defining features of selling information to competing buyers and show how information and competition interact in shaping the optimal mechanism. Both features distinguish the sale of information to competing firms from the sale of physical goods with externalities across buyers (e.g., network goods).

First, any buyer can always ignore (or indeed reverse) the seller's recommendation. The resulting *obedience* constraint limits the social planner's ability to reveal information to one buyer *exclusively*. Likewise, obedience disciplines the monopolist seller's ability to distort the allocation of information to maximize revenue at the expense of social welfare. Intuitively, the seller wants to distort the allocation of any buyer type with a negative Myersonian virtual value to minimize her payoff and reduce the information rents of higher types.

In our setting, this distortion corresponds to recommending the wrong action in every state. However, the buyer would not follow such a recommendation in any mechanism that does so too often. Therefore, the seller can do no better than to reveal "zero net information" to a low-value buyer, i.e., to probabilistically send the right and the wrong recommendation in a way that leaves the buyer indifferent over any course of action.

There are, of course, many such information structures (including entirely uninformative ones). However, the seller is not indifferent among them. Indeed, she can tailor the joint distribution of recommendations to maximize welfare while maintaining obedience on aggregate. The seller then prefers to reveal the correct state to all buyers when their types are sufficiently low. This approach relaxes obedience constraints and allows the seller to issue the correct recommendation to one or more buyers *exclusively* (and the wrong recommendation to the remaining buyers) when their types are sufficiently larger than their competitors'.

Second, providing information to a firm naturally imposes a negative externality on its competitors. In our setting, these negative externalities expand the profitability of selling information. Each buyer is willing to pay a positive price if either (a) she is strictly better off following the seller's recommendation or (b) her opponents do not receive the correct recommendation with probability one. As a result, the seller can charge a strictly positive price to some types with negative Myersonian virtual values, including some types whose obedience constraint binds, i.e., those who do not receive any valuable information themselves.

Leading Example Large digital platforms (e.g., Amazon, Facebook, and Google in the US, Alibaba and JD in China) collect an ever-increasing amount of information on their users' online behavior (e.g., browsing, shopping, social media interactions), which allows them to precisely estimate individual consumers' tastes for various products. Our leading application (fleshed out in Example 2) considers the interaction between a digital platform (information seller) and two or more merchants (information buyers). The merchants wish to leverage the platform's information advantage to offer a personalized product to each individual consumer.

The platform monetizes its proprietary information by selling *targeted* advertising slots (e.g., Facebook, Google) or sponsored marketplace listings (e.g., Amazon) to advertisers and retailers. Such practices amount to *indirect sales of information*: while the platform does not trade its consumers' data for payment (*direct sales*), it nonetheless creates value for merchants by allowing them to condition their strategies (in particular, which product to offer) on the consumers' information (e.g., their browsing or shopping history and third-party cookies). For the purposes of our model, direct transfers of information and indirect sales of information are, in fact, equivalent.[2]

Each merchant's expected volume of sales depends on two critical factors: (i) the degree of targeting, i.e., the precision of its advertising campaign, as measured by the ability to show the right product to each consumer; and (ii) the exclusivity of its campaign, i.e., the mismatch between its competitors' offers and the consumer's tastes. Merchants are willing to pay more for targeted campaigns and even more for exclusive access to targeted campaigns.

[2]In our approach, we further assume a platform such as Amazon has full commitment power to set information structures. Recent work by Koessler and Skreta (2022) analyzes the problem of information disclosure to multiple agents by a designer without commitment power.

However, the merchants' willingness to pay for an advertising campaign also depends on the profitability of making each sale, i.e., on their cost structure. As the latter is privately known to the merchant, the platform must elicit it through its choice of mechanism. Abstracting from the details and dynamics of online advertising auctions, the platform's problem reduces to designing a menu of (information structure, payment) pairs, each corresponding to an advertising campaign.

Related Literature This paper is primarily related to the literature on markets for information. In seminal work, Admati and Pfleiderer (1986, 1990) study the sale of information to traders who compete in financial markets. More recently, Babaioff et al. (2012) and Bergemann et al. (2018) study settings where a single buyer has private information about her beliefs over a payoff-relevant state. This problem is similar to ours insofar as the optimal mechanism can be represented through menus of information structures and associated prices, but there is a single buyer only.

Closer to our model, Rodríguez Olivera (2021) studies fully general mechanisms in a model with binary actions, states, and types. However, the buyers' types correspond to the realizations of a privately observed, exogenous signal about the state. Bimpikis et al. (2019) also study a setting similar to ours but consider mechanisms with a single option only—selling the true state distorted by Gaussian noise. Their problem then consists of finding the optimal covariance matrix of the noise and the associated prices. In particular, the covariance matrix is not designed as a function of the buyers' private types.[3]

Our work is also related to the recent literature on Bayesian persuasion and information design, e.g., Bergemann and Morris (2016), Kamenica (2019), Taneva (2019), and the references therein. Most papers in this literature view the the problem as a pure information design question as opposed to a mechanism design problem with transfers. In particular, these papers do not study the information structures that maximize the designer's revenue.

In a seminal paper, Jehiel et al. (1996) study an auction setting with multidimensional and interdependent valuations. Each buyer is privately informed about her valuation for a good and about the externality that she imposes on others. They show that the revenue-maximizing mechanism may involve not selling the good at all (when this is socially optimal) while charging positive payments to the buyers.[4] In further work, Jehiel and Moldovanu

[3] Xiang and Sarvary (2013) study information sellers who offer selling *exogenous* information structures about a binary state to buyers with *known* types who compete in a game with binary actions. Kastl et al. (2018) study the sale of cost information to a large number of perfectly competitive firms, each one facing a privately informed manager. Bounie et al. (2021) consider two sellers who acquire information about a consumer's location along a Hotelling line.

[4] Jehiel et al. (1999) study the simpler problem where each buyer knows the externality imposed on her by others receiving an object. Under appropriate symmetry assumptions, the problem reduces to a

(2000) restrict attention to the second-price auction and study more general externalities in the game played by the buyers. See Jehiel and Moldovanu (2006) for an exhaustive survey of the literature on mechanism design with externalities.

Our analysis is also very closely related to the model of data auctions with externalities in Agarwal et al. (2020). In their paper, the externalities resulting from the allocation of information are *intrinsic* to the buyers—the negative marginal effect of a competing buyer acquiring information is part of each buyer's private type. Finally, recent work by Kang (2022) and Pai and Strack (2022) explores a mechanism-design approach to the taxation of goods with externalities.

Relative to all these papers, our analysis highlights the differences between selling information and traditional goods in markets with *endogenous* downstream externalities. In contrast to the sale of physical goods, the allocation of information is both more flexible and more constrained. On the one hand, the seller has the flexibility to design any statistical experiment for each profile of buyer types. On the other hand, information is an input into the buyers' strategic decision problem (the "downstream game") that the seller does not control. As such, the sale of information introduces both obedience constraints, which are new to this literature, and tighter participation constraints.

2 Model

We consider n data buyers who compete in a downstream game of incomplete information. A monopolist data seller observes a payoff-relevant state variable and sells informative signals to the data buyers.

Notation For a tuple of sets $(\mathcal{S}_i)_{i\in[n]}$, we write $\mathcal{S} = \prod_{i=1}^n \mathcal{S}_i$ and $\mathcal{S}_{-i} = \prod_{j\neq i} \mathcal{S}_j$. Similarly for $s \in \mathcal{S}$, s_i (resp. s_{-i}) denotes the projection of s on $\mathcal{S}_i$ (resp. $\mathcal{S}_{-i}$). Finally $\Delta(\mathcal{S})$ denotes the set of probability distributions over $\mathcal{S}$.

Downstream Game We consider a downstream game of incomplete information between n buyers parametrized by an unknown parameter θ (the *state of the world*). We denote by Θ the set of all possible states. Each buyer $i \in [n]$ is described by a set of *types* $\mathcal{V}_i$, a set of

one-dimensional mechanism-design problem. In this vein, Ostrizek and Sartori (2021) study a screening model with externalities where each buyer's type affects both her valuation (e.g., for a network good) and the influence her actions (e.g., consumption) impose on other buyers. Their analysis focuses on the countervailing impact of payoff-types and influence functions.

actions $\mathcal{A}_i$, and a utility function

$$u_i : \mathcal{A} \times \Theta \times \mathcal{V} \to \mathbb{R}.$$

Information Structures The monopolist information seller chooses a set of signals $\mathcal{S}$, and a message (bid) space $\mathcal{B}$ to design a communication rule $\sigma : \Theta \times \mathcal{B} \to \Delta(\mathcal{S})$ and payment functions $p = (p_1, \dots, p_n) \in (\mathbb{R}^{\mathcal{B}}_{\geq 0})^n$. Given a vector of bids $b \in \mathcal{B}$ and state $\theta \in \Theta$, we write $\sigma(\,\cdot\,; \theta, b) : \mathcal{S} \to [0, 1]$ for the corresponding probability distribution over $\mathcal{S}$.

The buyers' utility functions $(u_i)_{i\in[n]}$, the mechanism (σ, p), as well as the joint distribution of the random variables $(\theta, V) \in \Theta \times \mathcal{V}$ are commonly known at the onset of the game.

Timing The interaction between the information seller and the buyers, and among the buyers in the downstream game, takes place as follows:

1. Each buyer $i \in [n]$ observes their type V_i and the seller observes the state θ.
2. Each buyer reports a message B_i to the information seller, where B_i is a V_i-measurable random variable in $\mathcal{B}_i$.
3. The information seller generates signals $S \in \mathcal{S}$ distributed as $\sigma(\theta, B)$ and reveals S_i to each buyer $i \in [n]$ in exchange for payment $p_i(B_i)$.
4. Each buyer i chooses an action $A_i \in \mathcal{A}_i$ that is (V_i, S_i)-measurable and obtains a total utility of $u_i(A; \theta, V) - p_i(B_i)$.

Remark 1. In the above formulation, the payment p_i of buyer i only depends on their own type report B_i. This is a departure from many mechanism design papers, where the payments are usually defined on the entire vector of type reports B. In an information design context, we could even consider more general payments that also depend on the state θ and the designer's signal S_i. As it happens, all these formulations reduce without loss of generality to the simple one above in which p_i only depends on B_i as we now explain.

Consider a general payment of the form $p_i(B, \theta, R)$, where R is the seller's sampling randomness, independent of all other variables. This formulation allows for randomized payments and subsumes in particular the case of payments depending on action recommendations since any (randomized) action recommendation can be written in the form $A_i = f_i(B, \theta, R)$ for some measurable function f_i. The key observation is that, because utilities are quasilinear, the seller's constraints (truthfulness, obedience, and individual rationality, cf. Section 3.1) only depend on p_i through the *interim* payment $\tilde{p}_i(v_i) = \mathbb{E}[p_i(B, \theta, R) \mid B_i]$. In other words, given any feasible mechanism with a general payment of the form $p_i(B, \theta, R)$, the mechanism in which we replace the payment with $p_i'(B_i) = \mathbb{E}[p_i(B,\theta,R) \mid B_i]$ is equivalent from the perspective of player i. It will thus satisfy the same constraints and lead to the same expected revenue. We therefore adopt the simpler form $p_i(B_i)$ in the rest of the paper for ease of notation.

The above formulation reduces the problem of information sale to a joint mechanism and information design problem. In order to obtain closed-form characterizations of the welfare- and revenue-optimal mechanisms in Section 4, we will restrict our attention to a specific case of this problem that we describe next. However, the characterizations of the constraints faced by the seller, which we study in Section 3, hold fore more general games.

Binary Game with Symmetric Additive Payoffs We consider a downstream game with n buyers, two states of the world, $\Theta = \{0,1\}$ and two actions for each buyer, $\mathcal{A}_i = \{0,1\}$ for $i \in [n]$. The utility of buyer $i \in [n]$ is given by

$$u_i(a;\theta,v) = v_i \cdot \pi_i(a;\theta).$$

In other words, the buyers have private *payoff types* that capture their marginal valuation for the downstream outcomes and reveal nothing about the state of the world. We further assume that the random variables $(\theta, v_1, \ldots, v_n)$ are drawn from a mutually independent prior. In this class of games, we assume the *downstream payoff* π_i of buyer $i \in [n]$ is given by

$$\pi_i(a;\theta) = \mathbf{1}\{a_i = \theta\} - \frac{\alpha}{n-1} \sum_{j \neq i} \mathbf{1}\{a_j = \theta\}. \tag{1}$$

Thus, in each state of the world, it is a dominant strategy for each buyer to play the action matching that state, resulting in a payoff gain of 1. A buyer additionally incurs a payoff loss of $\alpha/(n-1)$ whenever one of their competitors plays the action matching the state.[5] The externalities are thus normalized in such a way that α parametrizes the maximum externality that can be induced on a given buyer by the other buyers. Finally, given our focus on competitive environments, we assume that $\alpha \geq 0$.[6]

Example 2 (Binary Product Choice). The binary game described above can be seen as a stylized formulation of the motivating example presented in the introduction, with the state

[5]In the binary game, "choosing the correct action" is analogous to "being awarded the object" in an auction with externalities. In that case, the function π_i corresponds to the value of a given allocation for any buyer i, which is then scaled by the buyer's type v_i. Throughout the paper, and especially in Section 5, we will discuss significant differences between the allocation of physical and information goods.

[6]The case of positive externalities ($\alpha < 0$) is a straightforward extension of our analysis and is discussed in Section 6.1.

$\Theta = \{0, 1\}$ representing an individual consumer's preferences (unknown to the merchants). In this context, the merchants' goal is to match their product to the consumer's preferences. Each merchant is privately informed of its marginal profitability v_i in the downstream market. When there are only two merchants, we can write the payoffs (1) for each action profile $a \in \{0, 1\}^2$ in each state of the world θ as

	0	1
0	$1-\alpha, 1-\alpha$	$1, -\alpha$
1	$-\alpha, 1$	$0, \quad 0$

$\theta = 0$

	0	1
0	$0, \quad 0$	$-\alpha, 1$
1	$1, -\alpha$	$1-\alpha, 1-\alpha$

$\theta = 1$

The parameter $\alpha \geq 0$ captures the intensity of the competition between the merchants. A special case of this game occurs when $\alpha = 1$ in which case we have a zero-sum game. For $\alpha > 1$, the negative externalities outweigh the direct effect of choosing the correct action: the game turns into a prisoners' dilemma. Finally, each merchant is privately informed about its unit profit margin, i.e., merchant i's total payoff (gross of any payments to the seller) is given by $v_i \cdot \pi_i$.

3 Incentive Compatibility and Participation

3.1 Definitions

Incentive Compatibility In the game described in Section 2, each data buyer makes two strategic decisions: (i) report a message $B_i \in \mathcal{B}_i$ after observing their private type V_i, and (ii) take an action A_i in the downstream game after receiving signal $S_i \in \mathcal{S}_i$ from the seller.

By the revelation principle for dynamic games, see Myerson (1991, Section 6.3), it is without loss of generality to assume that the seller's set of signals $\mathcal{S}_i$ is equal to the set of actions $\mathcal{A}_i$, and that the buyers' reports lie in their own type space $\mathcal{V}_i$ instead of a general message space $\mathcal{B}_i$, as long as we consider incentive compatible mechanisms.

In any such mechanism, the seller recommends to the buyer an action to take in the downstream game. Henceforth, we therefore denote the seller's recommendation by A_i, and then buyer's choice of action by a_i. Incentive compatibility (below) requires each buyer to both report her true type and to follow the seller's recommendation.

Definition 1 (Incentive Compatibility)**.** A mechanism (σ, p) is *incentive compatible* if, for each $(v_i, v_i') \in \mathcal{V}_i^2$ and for each deviation function $\delta : \mathcal{A}_i \to \mathcal{A}_i$,

$$\mathbb{E}\big[u_i(A_i, A_{-i}; \theta, V) - p_i(v_i) \mid V_i = v_i, B_i = v_i\big] \geq \\ \mathbb{E}\big[u_i(\delta(A_i), A_{-i}; \theta, V) - p_i(v_i') \mid V_i = v_i, B_i = v_i'\big],$$

where A is distributed as $\sigma(\theta, B_i, V_{-i})$. The deviation function δ maps the seller's recommended action into the buyer's chosen action.

This definition of incentive compatibility is closely related to the one of Bergemann and Morris (2019, Section 3.1) in the context of information design with elicitation (but no transfers). In particular, Definition 1 requires the mechanism to be robust to *double deviations* in which the data buyer both misreports their private type and deviate from the seller's recommendation. This implies that the mechanism is both *truthful* and *obedient* as defined next.

Definition 2 (Obedience)**.** A mechanism (σ, p) is *obedient* if buyers have no incentive to deviate from the action recommendation of the seller assuming everyone reports their type truthfully. Formally, for each $\delta : \mathcal{A}_i \to \mathcal{A}_i$ and $v_i \in \mathcal{V}_i$,

$$\mathbb{E}\big[u_i(A_i, A_{-i}; \theta, V) \mid V_i = v_i\big] \geq \mathbb{E}\big[u_i(\delta(A_i), A_{-i}; \theta, V) \mid V_i = v_i\big]$$

where A is distributed as $\sigma(\theta, V)$—in particular, data buyer i's report is truthful.

Equivalently one can write obedience as: for each $(a_i, a_i') \in \mathcal{A}_i^2$ and $v_i \in \mathcal{V}_i$,

$$\mathbb{E}\big[u_i(a_i, A_{-i}; \theta, V) \mid V_i = v_i, A_i = a_i\big] \geq \mathbb{E}\big[u_i(a_i', A_{-i}; \theta, V) \mid V_i = v_i, A_i = a_i\big],$$

where A is distributed as $\sigma(\theta, V)$.

The first expression shows data buyer i's strategic behavior before receiving the action recommendation when she intends to report her type in the first stage of the game. At this stage, the buyer's strategy specifies a course of action following any action recommendation from the seller. Obedience requires that no deviations $\delta : \mathcal{A}_i \to \mathcal{A}_i$ are more profitable than obedience, i.e., the identity mapping $id : \mathcal{A}_i \to \mathcal{A}_i$.

The second expression shows data buyer i's strategic behavior after receiving the action recommendation at the second stage, and expresses that no other action results in a better expected utility. As mentioned before, these two are equivalent.

The second expression (which assumes every buyer reports her type truthfully) shows that obedience is only a property of the downstream game and of the recommendation rule σ, which thus correlates the actions of the data buyers. The distribution of actions resulting from an obedient recommendation rule in a game of incomplete information is a *Bayes correlated equilibrium* as defined and studied in Bergemann and Morris (2016, 2019).

Definition 3 (Truthfulness)**.** A mechanism is *truthful* if buyers have no incentive to misreport their type, assuming that everyone follows the seller's recommendations in the downstream game. Formally, for each $(v_i, v_i') \in \mathcal{V}_i^2$,

$$\mathbb{E}\big[u_i(A;\theta,V) - p_i(v_i) \mid V_i = v_i, B_i = v_i\big] \geq \mathbb{E}\big[u_i(A;\theta,V) - p_i(v_i') \mid V_i = v_i, B_i = v_i'\big]$$

where A is distributed as $\sigma(\theta, B_i, V_{-i})$.

Incentive compatibility implies both obedience and truthfulness, but the converse is not true in general. In Section 3.2, however, we show that with independent, private payoff types and linear valuations, incentive compatibility is equivalent to obedience and truthfulness.

Participation The data buyers engage in downstream competition even when they acquire no information from the seller. Thus, a complete description of the mechanism must specify the recommendations sent to the participating buyers when one ore more buyers choose not to participate in the mechanism. In that case, the data seller's recommendations to their competitors affect the non-participating buyers' utilities.

We define each buyer's bid space to be their type space $\mathcal{V}_i$ augmented with the special symbol '$\perp$', representing the decision not to participate. Writing $\mathcal{B}_i \coloneqq \mathcal{V}_i \cup \{\perp\}$ for the bid space of buyer $i \in [n]$ and $\mathcal{B} \coloneqq \prod_{i=1}^n \mathcal{B}_i$, the communication rule is now a function $\sigma : \Theta \times \mathcal{B} \to \Delta(\mathcal{A})$ with the constraint that it only sends a recommendation to the participating buyers. In other words,

$$\forall\theta \in \Theta,\ \forall b \in \mathcal{B},\ \sigma(\theta, b) \in \Delta(\textstyle\prod_{i:b_i \neq \perp} \mathcal{A}_i).$$

Similarly, the payment function $p_i : \Theta \times \mathcal{B} \to \mathbb{R}_{\geq 0}$ of buyer $i \in [n]$ satisfies the constraint that $p_i(\theta, b) = 0$ whenever $b_i = \perp$.

For each buyer $i \in [n]$, σ induces a communication rule $\sigma_i^o : \Theta \times \mathcal{V}_{-i} \to \Delta(\mathcal{A}_{-i})$ on the remaining buyers when buyer i chooses not to participate. This induced communication rule is given by,

$$\forall\theta \in \Theta,\ \forall v_{-i} \in \mathcal{V}_{-i},\ \sigma_i^o(\theta, v_{-i}) \coloneqq \sigma(\theta, \perp, v_{-i}).$$

This communication rule determines the *outside option* available to non-participating buyers: in any equilibrium where every buyer participates, any deviating buyer i chooses her action in the downstream game to be the best response to σ_i^o, resulting in the reservation utility

$$\max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[u_i(a_i, A_{-i}; \theta, V) \mid V_i = v_i\big] ,$$

where A_{-i} is distributed according to $\sigma_i^o(\theta, V_{-i})$. This is in marked contrast with a monopoly without externalities, in which a non-participating buyer simply receives no allocation, resulting in a vanishing reservation utility. It is also richer than in markets for physical goods with externalities, where a non-participating buyer has no available actions to choose from. We can now state the participation constraint.

Definition 4 (Individual Rationality)**.** The mechanism (σ, p) is individually rational each for each buyer $i \in [n]$,

$$\mathbb{E}\big[u_i(A; \theta, V) - p_i(V_i) \mid V_i = v_i\big] \geq \max_{a_i \in \mathcal{A}_i} \mathbb{E}\big[u_i(a_i, A_{-i}; \theta, V) \mid V_i = v_i\big] , \tag{2}$$

where A_{-i} is distributed according to $\sigma_i^o(\theta, V_{-i})$.

Intuitively, it is always in the seller's interest to relax this constraint as much as possible by selecting the outside communication rule σ_i^o that minimizes the right hand side in (2). In other words, the seller "punishes" a non-participating buyer by sending optimal recommendations to the remaining buyers so as to maximize the externalities induced on the deviating buyer. The specific way to achieve this depends on the downstream game and will be made explicit in Section 4.2.

Remark 3. The restriction to individually rational mechanisms is without loss of generality. Indeed, consider a mechanism for which some types do not participate at equilibrium. If we modify this mechanism to send the uninformative recommendation—matching the most likely state under the prior—to all non-participating types, we obtain a new mechanism in which the agents now (weakly) prefer to participate and play the same actions as in the original mechanism. In other words, any equilibrium can also be obtained as an equilibrium of a different mechanism in which everyone participates. In fact, the mechanisms we construct in Section 4 show that the seller can take advantage of agents' participation by inducing them to be correct when it hurts their competitors the least, thereby resulting in outcomes that could not be achieved without full participation.

3.2 Characterizations

In Section 4, we shall focus on the binary game with symmetric additive payoffs described in Section 2 and solve for the welfare- and revenue optimal mechanisms subject to the incentive compatibility and participation constraints defined in the previous section. To this end, this section provides characterizations of these two constraints.

Incentive Compatibility We begin the analysis with a characterization of incentive compatibility (Definition 1). We first restrict the buyers' utility to be multiplicatively separable in their independent private types and the outcome of the downstream game. In other words, the utility of buyer i depends linearly on their own type v_i, which is a buyer's marginal valuation for their *downstream payoffs* π_i.

Assumption 1 (Linear Payoffs)**.** *The random variables* $(\theta, V_1, \ldots, V_n)$ *are mutually independent. For each* $i \in [n]$*,* V_i *is supported on the non-negative reals and the utility of buyer* i *is given by*

$$u_i(a; \theta, v) = v_i \cdot \pi_i(a; \theta)$$

for some downstream payoff function $\pi_i : \mathcal{A} \times \Theta \to \mathbb{R}$*.*

As discussed above, incentive compatibility rules out double deviations and implies both truthfulness and obedience. Proposition 1 shows that, under Assumption 1, the converse is true and incentive compatibility reduces to requiring truthfulness and obedience separately. In other words, double deviations are not profitable whenever a mechanism is immune to single deviations.

Proposition 1 (Incentive Compatibility Characterization)**.** *Under Assumption 1, a mechanism is incentive compatible whenever it is truthful and obedient.*

Proof. See Section A. □

Truthfulness In order to characterize truthful mechanisms, we follow the classical result of Myerson (1981), which we restate in Proposition 2 below using our notation. Let (σ, p) be a mechanism and define for buyer $i \in [n]$, the interim downstream payoff $\tilde{\pi}_i(V_i) := \mathbb{E}[\pi_i(A; \theta) \mid V_i]$. We then have the following familiar characterization result.

Proposition 2 (Truthfulness Characterization)**.** *The mechanism* (σ, p) *is truthful if and only if for each buyer* i*:*

1. *The interim downstream payoff* $\tilde{\pi}_i$ *is non-decreasing.*

2. *The payment* p_i *is given for* $v_i \in \mathcal{V}_i$ *by*

$$p_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + p_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds \,. \tag{3}$$

Proof. See Section A. □

Obedience Next, we turn to a characterization of obedience for the binary game with symmetric additive payoffs (1).[7] In this game, the dominant strategy for each buyer in the absence of any signal about θ is to play the action corresponding to the most likely state under the prior. By construction, this is the correct action with probability

$$\mathbb{P}[A_i = \theta \mid V_i] = \max_{k \in \Theta} \mathbb{P}[\theta = k] =: P_{\max}. \tag{4}$$

The characterization of obedience in Proposition 3 below requires that following the recommended action makes a buyer more likely to be correct than if choosing an action under the common prior.

Proposition 3 (Obedience Characterization)**.** *In the binary game with additive payoffs* (1) *under Assumption 1, a recommendation rule is obedient if and only if for each $i \in [n]$, it holds almost surely that*

$$\mathbb{P}[A_i = \theta \mid V_i] \geq P_{\max}.$$

Proof. See Section A. □

In our characterization of optimal mechanisms below, we exploit the strength of this result, i.e., that obedience is a property of the marginal distribution of actions recommended to buyer i. In other words, the designer can flexibly correlate the buyers' actions and state, provided each buyer is recommended the right action often enough *on average*.

4 Optimal Mechanisms

We now turn social welfare and revenue maximization. We show below that, for the binary downstream game with additive payoffs in Eq. (1), both objectives can be written as a weighted sum of the probabilities that the mechanism recommends the dominant strategy to each buyer (see Eq. (5) below). Hence, we first describe in Section 4.1 an optimal mechanism for a general class of objective functions of this form, which we then instantiate in Section 4.2 and Section 4.3 to derive mechanisms maximizing social welfare and revenue, respectively.

[7] Our characterization does not in fact require externalities to be additively decomposable and holds more generally for all models in which the externality incurred by a player is independent of their own action. Formally, this is the class of models for which the downstream payoff of player i can be written $\pi_i(a;\theta) = \mathbf{1}\{a_i = \theta\} - E_i(a_{-i};\theta)$ for some function E_i.

4.1 Optimal Mechanisms

We consider a general objective function of the form

$$W \coloneqq \mathbb{E}\left[\sum_{i=1}^{n} w_i(V)\mathbf{1}\{A_i = \theta\}\right] = \sum_{i=1}^{n} \mathbb{E}\Big[w_i(V)\mathbb{P}[A_i = \theta \mid V]\Big] \tag{5}$$

for weight functions $w_i : \mathcal{V} \to \mathbb{R}$.

Expression 5 and the characterization of obedience obtained in Proposition 3 suggest a convenient parametrization of the seller's problem in terms of the functions $h_i : \mathcal{V} \to [0,1]$ given by $h_i(V) \coloneqq \mathbb{P}[A_i = \theta \,|\, V]$ for each player $i \in [n]$. These functions can easily be expressed in terms of the recommendation rule σ. Indeed, we have almost surely

$$\begin{aligned}\mathbb{P}[A_i = \theta \mid V] &= \mathbb{E}[\mathbf{1}\{A_i = \theta\} \mid V] \\ &= \sum_{\substack{a\in\mathcal{A} \\ a_i=\theta}} \mathbb{E}[\mathbf{1}\{A_1 = a_1, \dots, A_n = a_n\} \mid V] = \sum_{\substack{a\in\mathcal{A} \\ a_i=\theta}} \mathbb{E}[\sigma(a; \theta, V) \mid V].\end{aligned}$$

Conversely, Lemma 4 below shows that it is possible to construct a recommendation rule which has h_i as its marginals. In other words, any choice of the marginal functions h_i can be "realized" by a recommendation rule. Hence, as long as designer's objective and the constraints on the recommendation rule can be expressed in terms of $h_i(V)$, we will directly optimize over these quantities. An optimal information structure σ in this class can then be obtained using Lemma 4.

Lemma 4 (Recommendation Rule from Marginals)**.** *Let h_i be measurable functions from $\mathcal{V}$ to $[0,1]$ for $i \in [n]$, then there exists a recommendation rule $\sigma : \Theta \times \mathcal{V} \to \Delta(\mathcal{A})$ such that almost surely, $\mathbb{P}[A_i = \theta \mid V] = h_i(V)$ for $i \in [n]$.*

Proof. See Section A □

We now describe a general recommendation rule that optimizes criteria of the form (5), which include social welfare and seller revenue, subject to the obedience constraints. Recall the definition of $P_{\max}$ given in (4).

Proposition 5 (Optimal Mechanism)**.** *For the binary game with symmetric additive payoffs, consider an objective W of the form (5) where for $i \in [n]$, $w_i : \mathcal{V} \to \mathbb{R}$ is a measurable function such that the random variable $w_i(v_i, V_{-i})$ is non-atomic for each $v_i \in \mathcal{V}_i$. For $i \in [n]$ let $t_i^\star : \mathcal{V}_i \to \mathbb{R}$ be such that for all $v_i \in \mathcal{V}_i$,*

$$\mathbb{P}\Big[w_i(v_i, V_{-i}) \geq t_i^*(v_i)\Big] = P_{\max}.$$

Then the deterministic recommendation rule given by

$$A_i = \theta \quad \textit{if and only if} \quad w_i(v) \geq \min\{0, t_i^\star(v_i)\}$$

for $i \in [n]$, maximizes W subject to obedience.

Proof. See Section A. □

In order to gain intuition into the characterization of optimal mechanisms, note that the objective function W in (5) and the obedience constraints are separable. In other words, the optimization problem reduces to solving separately for each $i \in [n]$ and $v_i \in \mathcal{V}_i$:

$$\begin{gathered} \max \mathbb{E}\big[w_i(v_i, V_{-i})h_i(v_i, V_{-i})\big] \\ \text{s.t. } \mathbb{E}[h_i(v_i, V_{-i})] \geq P_{\max}, \end{gathered}$$

where, as above, $h_i(v) = \mathbb{P}[A_i = \theta \mid V]$ is the "allocation of correct information" to buyer i and takes values in $[0, 1]$ by definition. In the absence of the obedience constraint, the optimal solution would simply be to choose $h_i(v) = \mathbf{1}\{w_i(v) \geq 0\}$. If this violates the obedience constraint, we must also allocate information to some types where $w_i(v) < 0$, but we want to do so where the weight function w_i is as large as possible. Hence we should consider the smallest possible superlevel set of w_i that guarantees that the constraint is satisfied. This set corresponds to the level $t_i^\star(v_i)$ defined in the proposition statement.

4.2 Welfare Maximization

We now leverage Proposition 5 to characterize the welfare-optimal mechanism in our environment. For the binary game with additive payoffs (1) under Assumption 1, we can write the expected social welfare as

$$\begin{aligned} W &= \sum_{i=1}^{n} \mathbb{E}\Bigg[V_i\Big(\mathbf{1}\{A_i = \theta\} - \frac{\alpha}{n-1}\sum_{j\neq i}\mathbf{1}\{A_j = \theta\}\Big)\Bigg] \\ &= \sum_{i=1}^{n} \mathbb{E}\Bigg[\Big(V_i - \frac{\alpha}{n-1}\sum_{j\neq i} V_j\Big)\mathbf{1}\{A_i = \theta\}\Bigg]. \end{aligned} \tag{6}$$

Using the characterization of obedience from Proposition 3, the problem of maximizing social welfare subject to obedience can be written

$$\max \sum_{i=1}^{n} \mathbb{E}\left[\left(V_i - \frac{\alpha}{n-1}\sum_{j\neq i} V_j\right)\mathbf{1}\{A_i = \theta\}\right]$$
$$\text{s.t. } \mathbb{P}[A_i = \theta \mid V_i] \geq P_{\max}, \text{ for } i \in [n] \text{ and a.s.}$$

which is of the form (5). We can thus apply Proposition 5 and obtain the following characterization of the welfare-maximizing (second best) mechanism.

Proposition 6 (Welfare Optimal Mechanism)**.** *Consider the binary game with additive payoffs* (1) *under Assumption 1. Further, assume that the buyers' types are identically distributed with absolutely continuous c.d.f. F and denote by $F^{(k)}$ the c.d.f. of the sum of k i.i.d. variables*[8] *distributed according to F.*

Define $v^\star \in \mathbb{R}$ such that

$$F^{(n-1)}(v^\star) \coloneqq \mathbb{P}\left[\sum_{j\neq i} V_j \leq v^\star\right] = P_{\max}.$$

and $\overline{\alpha} \coloneqq \frac{\alpha}{n-1}$. Then, the recommendation rule maximizing social welfare subject to obedience is the deterministic rule given by

$$A_i = \theta \quad \textit{if and only if} \quad \sum_{j\neq i} v_j \leq \max\{v^\star, v_i/\overline{\alpha}\}.$$

Proof. See Section A. □

Figure 1 gives a representation of the welfare-optimal recommendation rule from Proposition 6 in the two-buyer case. This recommendation can be conceptualized as the "superposition" of two recommendation rules, which we describe separately in Figure 2.

1. The first rule (Figure 2, left) recommends the correct action to buyer i if and only if buyer j's type satisfies $v_j \leq v^\star$. For this rule, the recommendation to buyer i is independent or their type and satisfies $\mathbb{P}[A_i = \theta \mid V_i] = F(v^\star) = P_{\max}$. In other words, the recommendation is correct as often as buyer i would be by deterministically playing the action matching the most likely state under the prior. This implies by the characterization of Proposition 3 that the mechanism is obedient. Consequently, this mechanism recommends the correct action to buyer i *just often enough* to ensure

[8] $F^{(k)}$ can be computed recursively with $F^{(1)} = F$ and $F^{(k+1)} = F^{(k)} * f$, where $*$ denotes the convolution product and f is the p.d.f. associated with F.

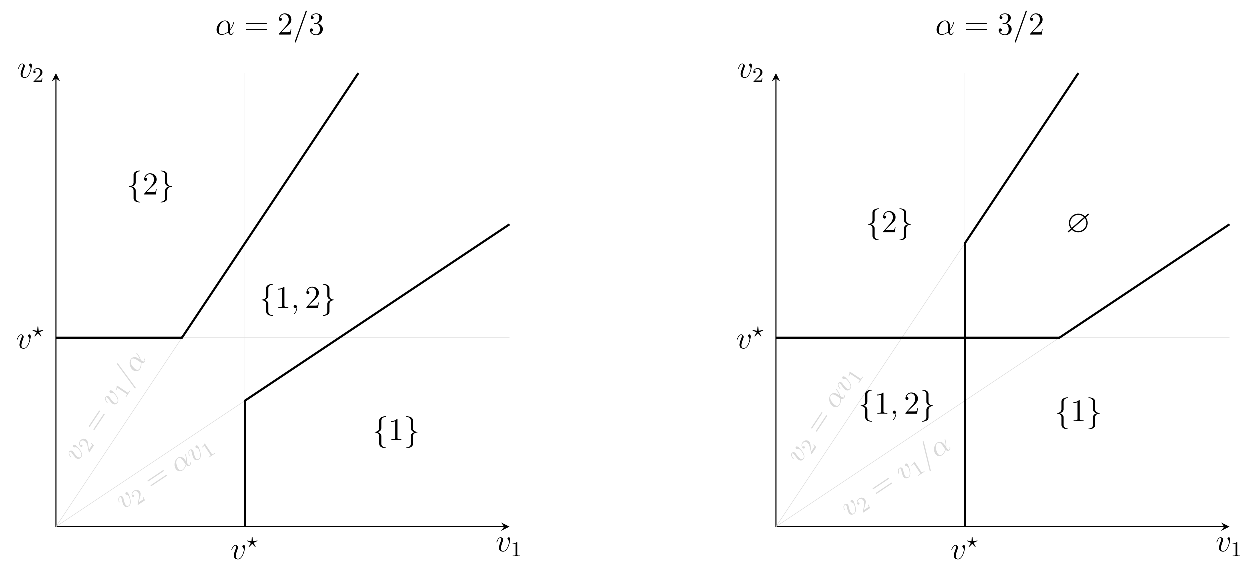

Figure 1: Welfare-maximizing recommendation rule from Proposition 6 with two buyers, for $\alpha = 2/3$ (left) and $\alpha = 3/2$ (right). The label in each region indicates the set of buyers who are recommended the correct action ($A_i = \theta$)—buyers in the complement set are recommended the wrong action ($A_i = 1 - \theta$). The two states are equally likely ex ante, so $v^\star = F^{-1}(1/2)$ is the median of the type distribution—chosen to be a standard exponential here).

obedience, and does so when buyer j's type is lowest, thus minimizing the induced externality $\alpha v_j \mathbf{1}\{A_i = \theta\}$. In summary, this mechanism ensures each buyer's obedience while minimizing the externality induced on the other buyer.[9] Note that the mechanism is obedient despite the action recommendations being deterministic in each region. This is because from the perspective of each buyer, conditional on their type, the recommendation they receive is still a random variable depending on the (unobserved) realization of the other buyer's type.

2. The second rule (Figure 2, center and right) recommends the correct action to buyer i if and only if her type satisfies $v_i \geq \alpha v_j$. This is simply the welfare-maximizing allocation (in the absence of the obedience constraint): a buyer is recommended the correct action if her value exceeds the externality she imposes on the other buyer. In particular, when buyer i's type is large enough compared to buyer j's type ($v_i/v_j \geq \max\{\alpha, 1/\alpha\}$), she is recommended the right action exclusively, hence maximizing her utility. In the intermediate region where types are close to each other, both buyers are recommended the same action. When $\alpha \leq 1$, the region is defined by $\alpha v_j \leq v_i \leq v_j/\alpha$ and the

[9] This also shows that the seller strictly benefits from the agents' participation (cf. Remark 3). Indeed, even for agents whose participation constraint is binding and who are thus receiving an action recommendation that is only correct with probability $P_{\max}$, the seller can control *when* the agent is correct *over the realizations of their competitor's type.*

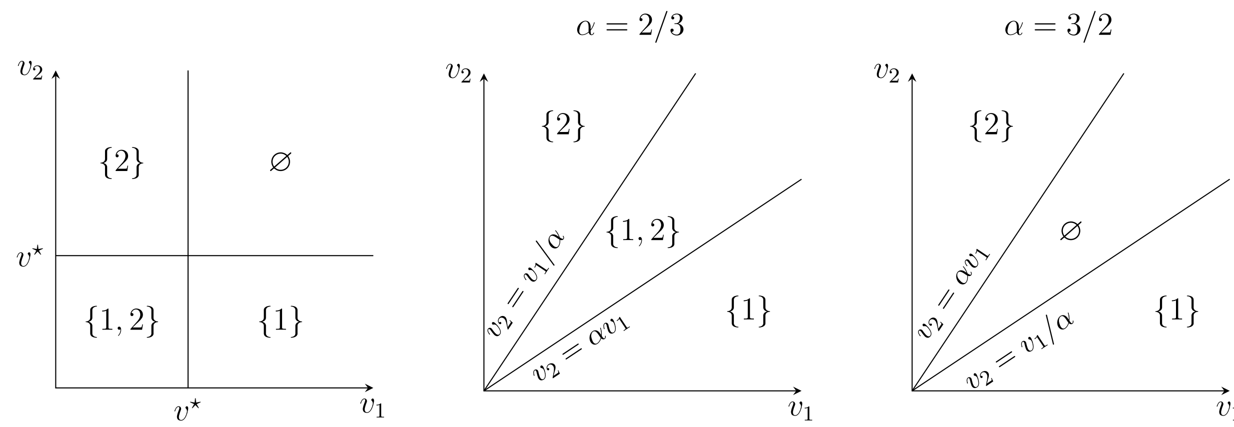

Figure 2: Building blocks for the welfare-maximizing mechanism of Proposition 6. Left: mechanism guaranteeing obedience at all types while minimizing externalities. Center and right: first-best mechanism (ignoring the obedience constraint) for $\alpha = 2/3$ and $\alpha = 3/2$.

efficient allocation recommends the correct action to both buyers. In contrast, when $\alpha > 1$, the region is defined by $v_j/\alpha \le v_i \le \alpha v_j$ and both buyers are recommended the wrong action. Indeed, the externalities are so significant in this case that the buyers face a prisoners' dilemma in each state. It is thus more efficient for the data seller to coordinate the buyers on the collaborative strategy in which both buyers pick the "wrong" action.

The optimal mechanism (Figure 1) combines both mechanisms by distorting the first best mechanism to guarantee that each buyer i receives the correct action when $v_j \le v^\star$. Distorting buyer i's recommendation is required, and hence obedience is binding, when $v_j \le v^\star$ and $v_i \le \alpha v_j$.

Finally, it is easy to verify that the second best mechanism is implementable, i.e., it satisfies the buyers' truth-telling constraints. Indeed, by Proposition 2 it suffices to verify that the interim downstream payoff is non-decreasing in the buyer's type.

Proposition 7 (Implementability of Second-Best Mechanism)**.** *For the deterministic mechanism of Proposition 6, and under the same assumptions, the interim expected payoff of buyer* $i \in [n]$, $\tilde{\pi}_i(V_i) \coloneqq \mathbb{E}\big[\pi_i(A;\theta) \mid V_i\big]$, *satisfies almost surely*

$$\tilde{\pi}_i(v_i) = \max\big\{F^{(n-1)}(v^\star), F^{(n-1)}(v_i/\overline{\alpha})\big\} - \overline{\alpha} \sum_{j \neq i} \mathbb{E}\Big[F^{(n-2)}\big(\max\{v^\star, V_j/\overline{\alpha}\} - v_i\big)\Big].$$

In particular, $\tilde{\pi}_i$ *is non-decreasing and the recommendation rule is therefore implementable.*

Proof. See Section A. □

Intuitively, a higher type is revealed the correct state more often by the social planner, which makes it possible to find transfers that would induce truthful reporting of the buyers' types. Of course these transfers do not correspond to a monopolist data seller's optimal choice. In the next section, we will see how a monopolist data seller modifies the second best mechanism to maximize the associated payments.

4.3 Revenue Maximization

Throughout this section, we further assume that the type distribution F is absolutely continuous with p.d.f. f and that the *virtual value function* $\phi : \mathcal{V}_i \to \mathbb{R}$ defined by

$$\phi(v) := v - \frac{1 - F(v)}{f(v)},$$

is non-decreasing, that is, F is *regular* in the sense of Myerson (1981).

We first show in Lemma 8 that maximizing the seller's expected revenue reduces to maximizing the virtual surplus, as in Myerson (1981).

Lemma 8 (Reduction to Virtual Surplus)**.** *Let σ be a communication rule for which the interim payoff $\tilde{\pi}_i$ is non-decreasing for each buyer $i \in [n]$. Denote by K the interim downstream payoff of a non-participating buyer*[10] *and assume that $\tilde{\pi}_i(\underline{v}) \geq K$. Then:*

1. *If p_i is a payment function that truthfully implements $\tilde{\pi}_i$ (i.e., that satisfies (3) by Proposition 2), then (σ, p) is individually rational if and only if it is individually rational for the lowest type, that is, $p_i(\underline{v}) \leq \underline{v} \cdot (\tilde{\pi}_i(\underline{v}) - K)$.*

2. *Among the payment functions p_i implementing $\tilde{\pi}_i$ in a truthful and individually rational manner, the revenue-maximizing one is given by*

$$p_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot K - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds\,. \tag{7}$$

For this payment function, the seller's revenue is $R = \sum_{i=1}^{n} \mathbb{E}\big[\phi(V_i)\tilde{\pi}_i(V_i)\big] - n\underline{v} \cdot K$.

Proof. See Section A. □

We thus focus on maximizing the virtual surplus $R^\dagger := \sum_{i \in \{1,2\}} \mathbb{E}\big[\phi(V_i)\tilde{\pi}_i(V_i)\big]$ subject to

[10] Using the notations of Definition 4, if σ_i^o denotes the recommendation rule used with the remaining buyers when buyer i does not participate, then we have $K = \mathbb{E}[\pi_i(a^*, A_{-i}; \theta)]$, where A_{-i} is distributed according to $\sigma_i^o(\theta, V_{-i})$ and a^* is the action matching the most likely state under the prior.

obedience and truthfulness. For the binary game (1), we write the virtual surplus as

$$R^\dagger = \sum_{i=1}^{n} \mathbb{E}\left[\Big(\phi(V_i) - \frac{\alpha}{n-1}\sum_{j\neq i}\phi(V_j)\Big)\mathbf{1}\{A_i = \theta\}\right].$$

This objective function is of the form (5) and we can thus apply Proposition 5 to characterize the communication rule maximizing virtual surplus subject to obedience. Then, we verify that the corresponding expected downstream payoff, $\tilde{\pi}_i$, is non-decreasing, implying that the mechanism is implementable in a truthful and individually rational manner using the payments given by (7).

Proposition 9 (Revenue Optimal Mechanism)**.** *Consider the binary game with additive payoffs (1) under Assumption 1. Further assume that the buyers' types are identically distributed with absolutely continuous c.d.f. F. Denote by F_ϕ the c.d.f.*[11] *of $\phi(V_i)$ where V_i is distributed according to F and by $F_\phi^{(k)}$ the c.d.f. of the sum of k i.i.d. variables distributed according to F_ϕ.*

Define $v^\star$ such that $F_\phi^{(n-1)}\big(\phi(v^\star)\big) = P_{\max}$ and $\overline{\alpha} \coloneqq \alpha/(n-1)$. Then, the recommendation rule maximizing virtual surplus subject to obedience is the deterministic rule given by

$$A_i = \theta \quad \textit{if and only if} \quad \sum_{j\neq i}\phi(v_j) \leq \max\{\phi(v^\star), \phi(v_i)/\overline{\alpha}\}.$$

Proof. The proof is identical to the one of Proposition 6 with $\phi(V_i)$ playing the role of V_i. It follows from an application of Proposition 5 with weight function $w_i(v) = \phi(v_i) - \overline{\alpha}\sum_{j\neq i}\phi(v_j)$. □

The functional form of the revenue-optimal mechanism in Proposition 9 is analogous to that of the welfare-optimal mechanism in Proposition 6, after replacing the buyers' types with their virtual types. Figure 3 shows the resulting recommendation rule for $n = 2$ buyers both when $\alpha < 1$ and $\alpha > 1$. Again, this recommendation can be understood as the superposition of two recommendation rules:

1. The first rule recommends the correct action to buyer i if and only if the virtual type of the other buyer satisfies $\phi(v_j) \leq \phi(v^\star)$, or equivalently since F is regular, $v_j \leq v^\star$. This is exactly the same mechanism as was shown in Figure 2 (left) guaranteeing the obedience of buyer i.

[11] When F is a regular distribution, the virtual value function ϕ is invertible and the c.d.f. F_ϕ can be computed as $F \circ \phi^{-1}$.

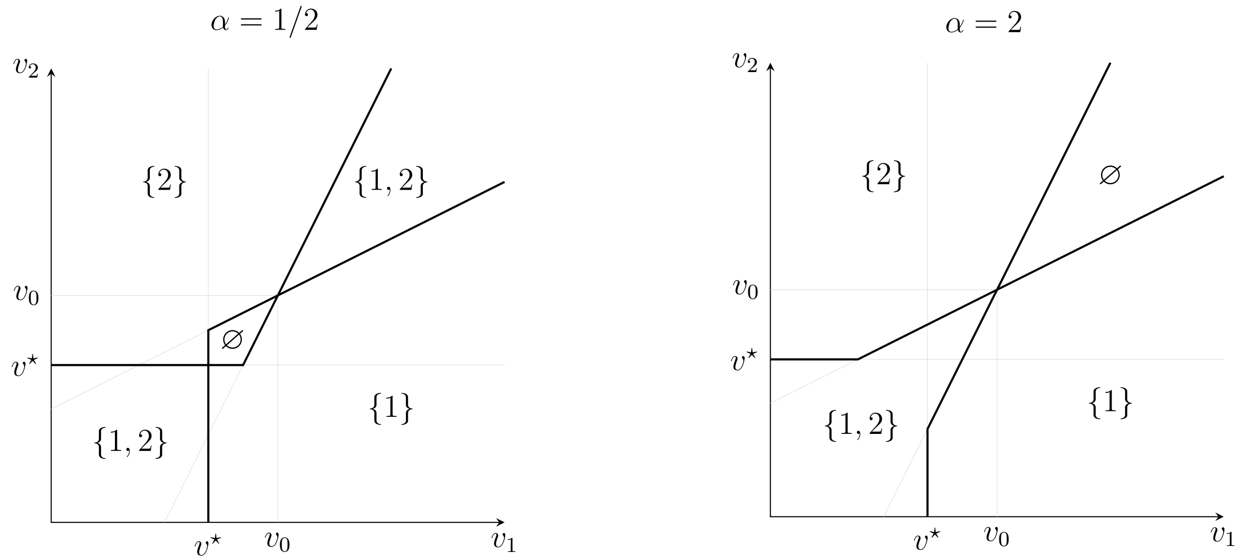

Figure 3: Revenue-maximizing recommendation rule from Proposition 9 for $\alpha = 1/2$ (left) and $\alpha = 2$ (right). Types are distributed exponentially, so that $\phi(v) = v - 1$ and $v_0 = \phi^{-1}(0) = 1$. The prior on θ is symmetric ($P_{\max} = 1/2$), hence $v^\star = F^{-1}(1/2) = \ln 2 < v_0$.

2. The second rule recommends the correct action to buyer i if and only if $\phi(v_i) \leq \phi(v_j)/\alpha$. In particular, when one virtual valuation is large compared to the other ($\phi(v_i)/\phi(v_j) \geq \max\{\alpha, 1/\alpha\}$), buyer i is recommended the correct action exclusively. However, because the functions $v \mapsto \phi^{-1}(\phi(v)/\alpha)$ and $v \mapsto \phi^{-1}(\alpha\phi(v))$ intersect at $v_0 := \phi^{-1}(0)$, the intermediate regime $\phi(v_i)/\phi(v_j) < \max\{\alpha, 1/\alpha\}$ now determines two regions in which both buyers are recommended the same action. When virtual valuations are positive (types greater than v_0), both buyers are recommended the correct action when $\alpha < 1$ and the wrong action when $\alpha > 1$. Indeed, in this latter case, the buyers face a prisoners' dilemma in which coordinating on the dominated "wrong" action results in higher payoffs. Naturally, the situation is reversed when virtual values are negative in the intermediate regime: both buyers receive the wrong action when $\alpha < 1$ and the correct one when $\alpha > 1$. This is shown in Figure 4 below.

The revenue-optimal mechanism resulting from the superposition of these two mechanisms depends both qualitatively and quantitatively on the relative positions of v_0 and $v^\star$. This in turn depends on the magnitude of the parameter $P_{\max}$ and is discussed in Section 5.2 below.

Proposition 10 below gives an expression for the expected downstream payoff $\tilde{\pi}_i$ of each buyer in the obedient mechanism described above. Note that due to ϕ being non-decreasing (since we assumed that the distribution F is regular), as buyer i increases their bid v_i, the recommendation rule in Proposition 9 recommends the correct action to i more often and to i's competitors less often. Both of these factors contribute to increasing buyer i's downstream payoff, which in turns implies that the interim payoff is non-decreasing as stated in the proposition. Consequently, the mechanism above is also *truthful (implementable)* and the payments are then given by Lemma 8. Given that these payments are decreasing as a function on K, the downstream payoff of a non-participating buyer, we must therefore design the outside option so as to minimize K. For the binary game with additive payoffs, the following proposition establishes that the optimal allocation when i does not participate recommends the correct action to the set $[n] \setminus \{i\}$ of all participating buyers.

Proposition 10 (Implementation of Optimal Mechanism)**.** *For the mechanism of Proposition 9 and under the same assumptions, the interim downstream payoff* $\tilde{\pi}_i$ *of buyer* $i \in [n]$ *is the non-decreasing function*

$$\tilde{\pi}_i(v_i) = \max\left\{F_\phi^{(n-1)}\big(\phi(v^\star)\big), F_\phi^{(n-1)}(\phi(v_i)/\overline{\alpha})\right\} \\ - \overline{\alpha}\sum_{j \neq i} \mathbb{E}\left[F_\phi^{(n-2)}\big(\max\{\phi(v^\star), \phi(V_j)/\overline{\alpha}\} - \phi(v_i)\big)\right],$$

and the revenue-maximizing mechanism is therefore implementable in a truthful manner.

In case of non-participation of buyer $i \in [n]$*, the recommendation rule minimizing their reservation utility recommends the correct action to the remaining buyers* ($A_j = \theta$ *for* $j \neq i$)*. For this outside option, the payments maximizing revenue subject to individual rationality and truthfulness are given by*

$$p_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s)ds + \underline{v}\alpha - \underline{v} \cdot P_{\max}.$$

Proof. See Section A. □

For the outside option in Proposition 10, the optimal strategy of a non-participating buyer is simply to play the action matching the most likely state under the prior, resulting in the buyer being correct with probability $P_{\max}$. Furthermore, the externality incurred by a non-participating buyer is $(n-1)\overline{\alpha} = \alpha$, because all participating buyers receive the correct action recommendation in this case. Hence, the reservation utility of a non-participating buyer is $P_{\max} - \alpha$: this is precisely the offset appearing in the expression for p_i, in Proposition 10 guaranteeing buyer i's participation.

We now remark on several properties of the optimal payments, which apply whenever $v^\star < v_0$, as in Figure 3.

1. Unlike in settings without externalities, merely having a negative virtual value does not imply a buyer receives no information. Even absent obedience constraints, the seller

knows that distorting one buyer's recommendation increases the surplus of the other buyer. Therefore, when $\alpha < 1$ both buyers receive the wrong recommendation only if both their virtual values are negative *and* they are sufficiently similar. Conversely, if both virtual values are negative but v_1 is sufficiently larger than v_2, then the seller prefers issuing the correct recommendation to buyer 1. Indeed, distorting the recommendation to buyer 1 would increase buyer 2's payoff, which has an even stronger negative impact on the seller's profits.

2. Some types of buyer i with a negative virtual valuation $v_i < v_0$, are nonetheless charged a positive payment. This occurs because these types are sufficiently high that their opponent j has an even lower type v_j with a significant probability, $F(v_i)$. In other words, the seller finds it optimal to reveal the correct state to buyer i with probability, $F(\phi^{-1}(\phi(v_i)/\alpha)) > F(v^\star)$. Buyer i then has a strict incentive to follow the seller's recommendation, i.e., her obedience constraint is slack.

3. Some types of buyer i such that $v^\star < v_i$, whose obedience constraint binds, still pay a strictly positive price. Because their obedience constraint is binding, these types derive no net utility from following the seller's recommendation. However, unlike types in $[0, v^\star]$ where the other data buyer always receives the right recommendation, these types' opponent is revealed the correct state with probability $1 - F(\phi^{-1}(\alpha\phi(v_i)))$. These types are strictly better off participating, and they can be charged a positive payment. Thus, the presence of negative externalities augments the profitability of selling information, as the seller charges positive payments in exchange for limiting the information available to each buyer's competitors.

5 Information and Competition

In this section, we discuss the impact of the environment facing the buyers on the optimal mechanisms presented in Section 4. This encompasses the seller's constraints, discussed in Section 5.1, the information structure and in particular the buyers' prior information discussed in Section 5.2, the competition structure in the downstream game as captured by the externality parameter α (Section 5.3), and the number of buyers (Section 5.4).

5.1 Benchmarks

In this section, we contrast our optimal mechanism to "benchmark" situations in which we relax either or both of our constraints: obedience and truthfulness.

1. With complete information (i.e., the seller knows v) and no obedience constraint, both the welfare- and revenue-optimal mechanism use the first-best recommendation rule shown in Figure 2 (center):

$$A_i = \theta \quad \text{if and only if} \quad \sum_{j \neq i} v_j \leq v_i / \overline{\alpha}. \tag{8}$$

Since there is no constraint on the payments in this case—other than incentivizing participation—the revenue optimal mechanism collects as a payment the difference between (i) the interim utility under allocation (8) and (ii) the reservation utility. This payment is given by,

$$p_i(v_i) = v_i\big(\tilde{\pi}_i(v_i) - P_{\max} + \alpha\big),$$

where the interim payoff $\tilde{\pi}_i(v_i)$ can be computed from (8) following the same steps as in Proposition 7

$$\tilde{\pi}_i(v_i) = F^{(n-1)}(v_i/\bar{\alpha}) - \bar{\alpha} \sum_{j \neq i} \mathbb{E}\big[F^{(n-2)}(V_j/\bar{\alpha} - v_i)\big]. \tag{9}$$

2. In the case of private information and no obedience constraints, our characterization of truthfulness in Proposition 2 requires the interim payoff to be non-decreasing. It is easy to verify that $\tilde{\pi}_i$ defined in (9) is non-decreasing, hence the recommendation rule (8) is truthful. Consequently, the welfare-optimal recommendation rule is identical in benchmarks (1) and (2). A payment that truthfully implements this rule is given in Proposition 2.

   Equivalently, as we will also discuss in Section 5.2, ignoring the obedience constraint amounts to formally setting $P_{\max} = 0$ in our mechanisms (in which case the obedience constraint becomes void). Using this equivalence, we immediately obtain that the revenue optimal recommendation with private information and no obedience is given by

$$A_i = \theta \quad \text{if and only if} \quad \sum_{j \neq i} \phi(v_j) \leq \phi(v_i)/\overline{\alpha}.$$

   with the associated payment described in Proposition 10, where the critical type $v^\star$ is now given by $\phi(v^\star) = \underline{v}$. This mechanism is the one depicted in Figure 6 (first row).

3. With complete information and obedience constraints, the optimal mechanisms will in fact be identical to the optimal mechanisms derived in Sections 4.2 and 4.3 respectively. This is because we derived these mechanisms by first ignoring the truthfulness constraint,

and then verifying that the interim payoffs were non-decreasing (this requires regular type distributions in the revenue case). In other words, truthfulness was obtained "for free" and this benchmark does not differ from the optimal mechanisms that take all the constraints into account.

5.2 Buyers' Prior Information

The seller's information augments the buyers' prior information and allows them to tailor their actions to the state of the world. But each buyer also has the option of playing the downstream game under their prior information only. Thus, each buyer's participation constraint is tighter when buyers are better informed, and the seller cannot extract all the buyers' surplus through transfers.

Furthermore, while the seller is unconstrained in her choice of experiments, the buyers retain the flexibility to choose their actions after observing the signals. These signals must then be sufficiently informative *relative to the buyers' prior* in order for the data buyers to follow them. In particular, the buyers can always ignore the recommendation altogether and choose the action that is optimal under the prior, or choose actions that respond to signals in a different way than the seller intended. Thus, not all distributions over action profiles in the downstream game are feasible for the seller, due to the buyers' obedience constraints.

In our binary setting, the seller's problem therefore depends critically on a scalar parameter: the informativeness of the buyers' prior beliefs, as captured by $P_{\max} := \max_{k \in \{0,1\}} \mathbb{P}[\theta = k]$. Indeed, $P_{\max}$ describes both novel aspects of our seller's problem: the buyer's reservation utility that corresponds to foregoing participation in the mechanism (in which case the seller fully reveals the state to the $n-1$ other buyers); and the buyer's option value of participating but ignoring the seller's recommendations. Formally, the obedience and participation constraints can be written as

$$\begin{aligned} \mathbb{P}[A_i = \theta \mid V_i = v_i] &\geq P_{\max}, \\ v_i \tilde{\pi}_i(v_i) - p_i(v_i) &\geq v_i \left(P_{\max} - \alpha\right). \end{aligned}$$

Figure 4 illustrates the revenue-maximizing mechanism for $n = 2$ buyers as we vary the parameter $P_{\max}$, for both $\alpha < 1$ and $\alpha > 1$. The top row describes the benchmark case where we artificially set $P_{\max} = 0$ in the constraints above. This corresponds to a setting akin to the sale of physical goods with externalities: there are no downstream actions for buyers to take (hence no obedience constraints), and each buyer's reservation utility consists of not receiving the good while their competitors all receive it with probability 1, so that $\tilde{\pi} = -\alpha$ (see also Section 5.1 (2)).

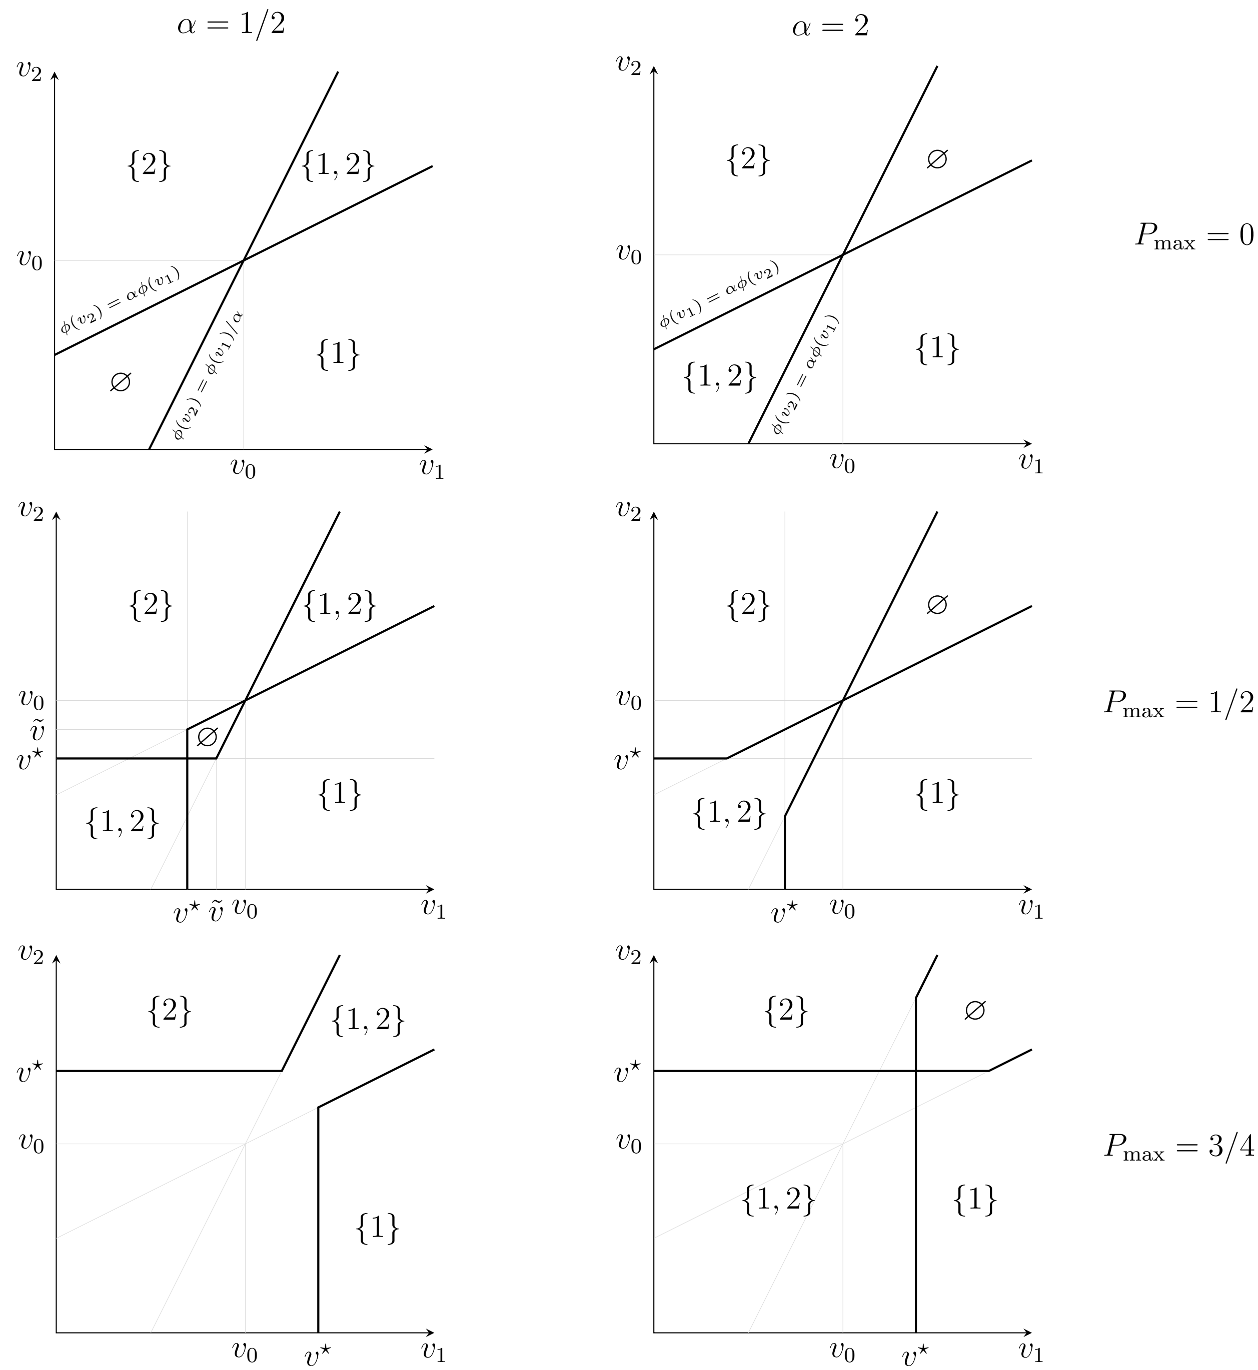

Figure 4: Revenue-maximizing recommendation rule from Proposition 9 for $\alpha = 1/2$ (left) and $\alpha = 2$ (right). Types are distributed exponentially, so that $\phi(v) = v - 1$ and $v_0 = \phi^{-1}(0) = 1$. The first row shows the first best mechanism. The second row is the second-best mechanism (subject to obedience) with a symmetric prior on θ, for which $v^\star = F^{-1}(1/2) = \ln 2 < v_0$. The third row is the second-best mechanism with an asymmetric prior ($p_{\max} = 3/4$), for which $v^\star = \ln 4 > v_0$.

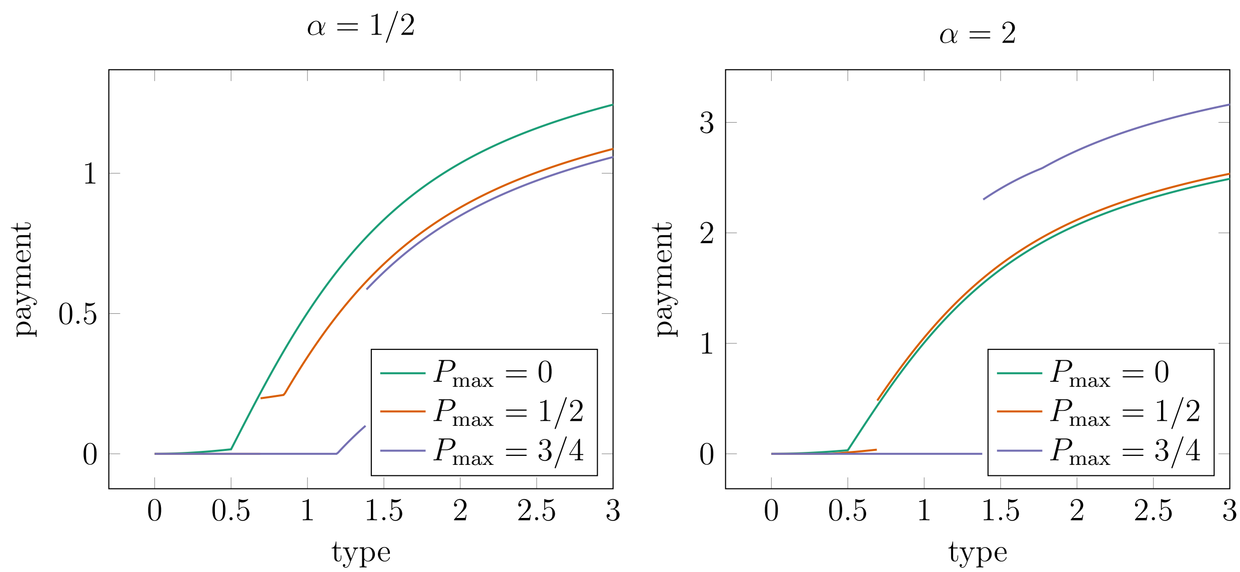

Figure 5: Payment as a function of a buyer's type, for different values of $P_{\max}$ with exponentially distributed types. Left panel: $\alpha = 1/2$. Right panel: $\alpha = 2$.

As $P_{\max}$ increases, as in the second row, the revenue-maximizing mechanism must assign the correct action to both buyers whenever *their competitor's* type is below the critical level $v^\star := F^{-1}(P_{\max})$. When $v^\star \geq v_0$ (as in the third row), the square where $v_1, v_2 \leq v^\star$, in which obedience requires recommending the optimal action to both buyers, fully contains the region with negative virtual valuations and the situation looks qualitatively the same as Figure 1.[12] In all cases obedience binds for buyer i for all types such that $\phi(v_i) \leq \alpha\phi(v^\star)$, or equivalently $v_i \leq \tilde{v} \coloneqq \phi^{-1}(\alpha\phi(v^\star))$.

Finally, Figure 5 shows the payments under the revenue-optimal mechanism for several values of $P_{\max}$. As previewed in Section 4.3, the model without obedience constraints ($P_{\max} = 0$) has positive payments for almost all types. In contrast, the precision of the buyers' prior information (captured by $P_{\max} \geq 1/2$) prevents the seller from charging any payments to a positive measure of types, despite the presence negative externalities. This is in sharp contrast with the sale of *final goods* with externalities, e.g., Jehiel et al. (1996).

5.3 Externality Parameter α

We now investigate the effect of the intensity of downstream competition on the revenue-optimal mechanism. Figure 6 compares two settings, where competition is fiercer in the left panel ($\alpha = 1/2$) than in the right panel ($\alpha = 1/4$).

[12] Depending on the type distribution, we may have $v^\star \geq v_0$ for any prior on the unknown state. For example, for types uniformly distributed over $[0, 1]$, $\phi(v) = v - 1$ and $v_0 = \phi^{-1}(0) = 1/2 \leq P_{\max} = v^\star$.

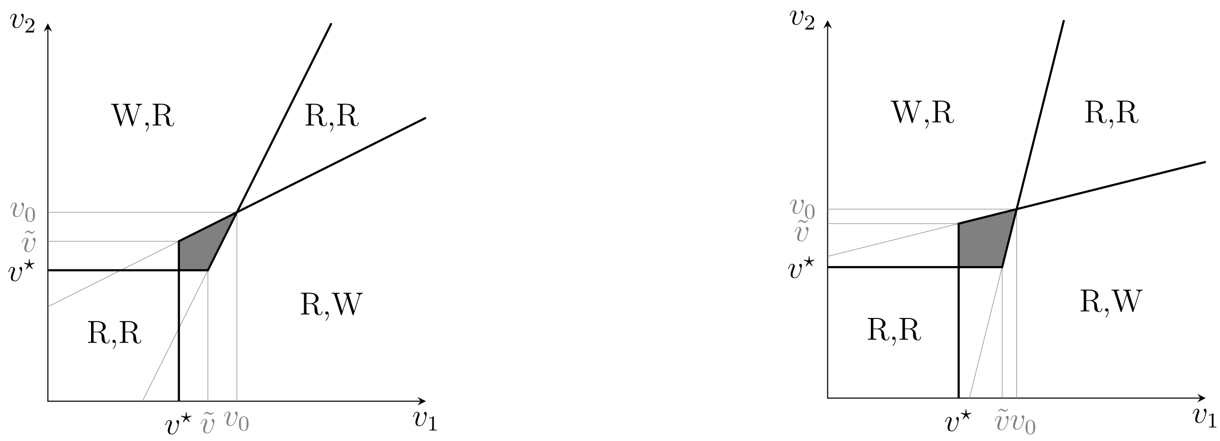

(a) $\alpha = 1/2$

(b) $\alpha = 1/4$

Figure 6: Comparison of the revenue-maximizing recommendation rules from Proposition 9 for two different values of α. The mechanism recommends the wrong action $(1-\theta)$ to both buyers in the gray regions. As in Figure 4, types are exponentially distributed and states are equally likely: the larger the value of α, the more competitive the downstream game.

Reducing the intensity of competition reduces the value of exclusive sales of information (i.e., recommending the right action to one buyer only) in the first best: at one extreme, if buyers imposed no externalities on each other, the seller would recommend the right action to any buyer with a positive virtual value. In particular, in an unconstrained revenue problem, the seller would recommend the wrong action to both buyers more often when competition is weaker. However, this recommendation profile would violate obedience, which requires the seller to recommend the right action to both buyers when both their types are smaller than $v^\star$. As $v^\star$ is independent of α, the right panel shows how the seller uses exclusive sales as a second-best policy under obedience constraints more often as the competition weakens.

In Figure 7, $v^\star < v_0$, downstream payoffs are non-decreasing as expected. When competition is more intense i.e. $\alpha = 1/2$ the expected downstream payoff, $\tilde{\pi}_i$, is smaller while payments are larger. This shows that the designer uses the competition between the firms as a tool to extract more surplus in exchange for the provision of exclusive information.

In Figure 8, $v_0 < v^\star$, the plots follow a similar pattern as Figure 7 except for types which are less than $v^\star$. For these types, fiercer competition results in lower payments. The increase in competition does not provide data buyers with exclusive information, and at the same time their expected downstream payoff decreases. Therefore, they do not have any incentive to pay more. However, Figure 9 suggests that the overall expected payment $\mathbb{E}\left[p_i(V_i)\right]$ increases with α.

Figure 9 shows that welfare decreases as competition increases. By increasing α the entries

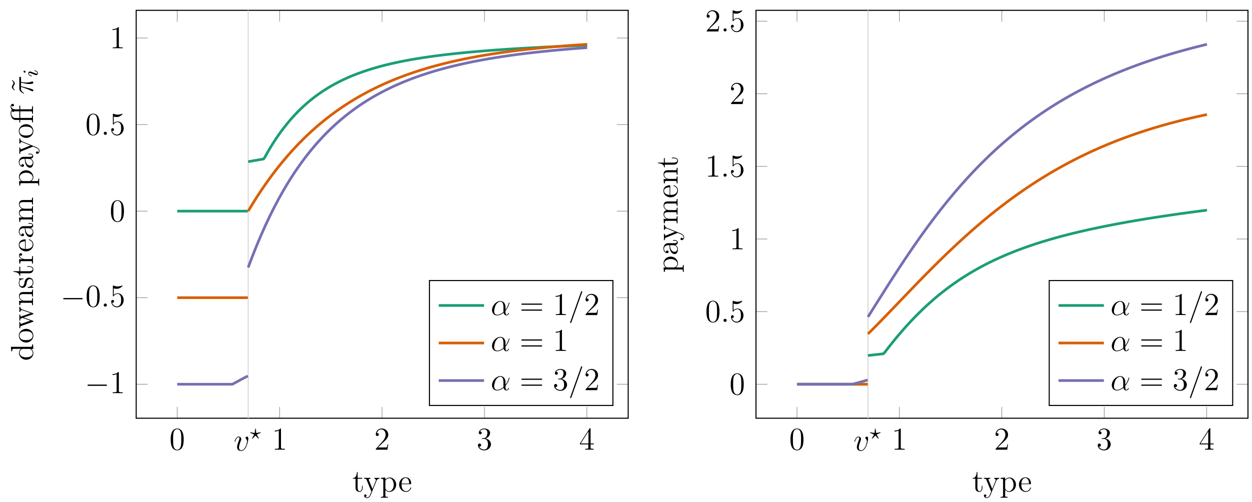

Figure 7: Interim downstream payoff and payment as functions of a buyer's type in the optimal mechanism for different values of α. Types are exponentially distributed and the prior on the unknown state is uniform ($P_{\max} = 1/2$). There is a discontinuity at $v^\star = \ln 2 < v_0$ and a singularity at $\tilde{v}_\alpha = \phi^{-1}\big(\alpha\phi(v^\star)\big)$. The minimum expected downstream payoff, $\tilde{\pi}_i(\underline{v})$, is $P_{\max} - \alpha$.

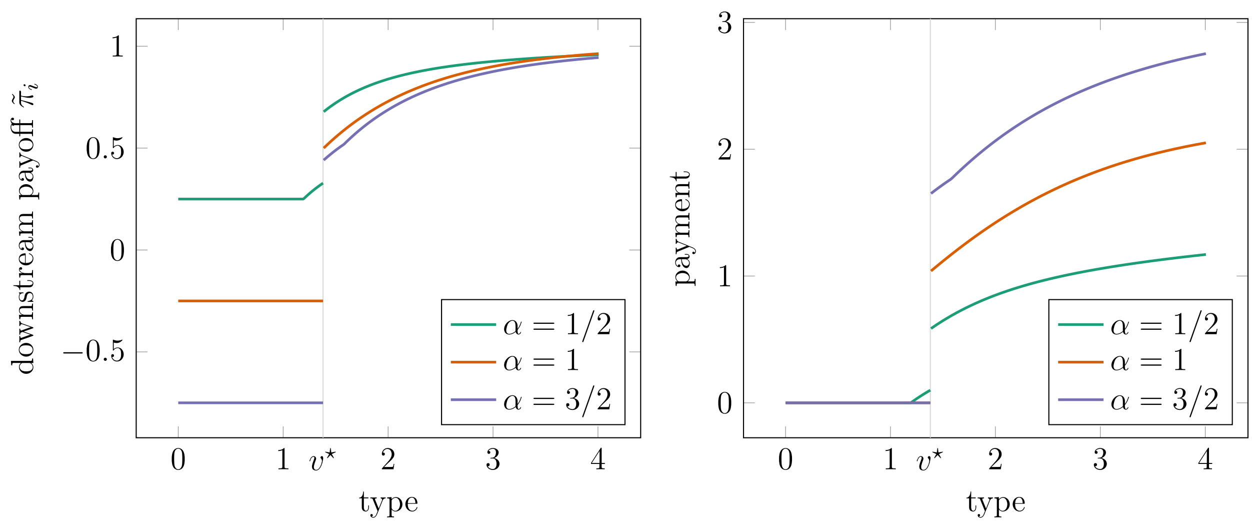

Figure 8: Interim downstream payoff and payment. Same parameters as Figure 7 except that now $P_{\max} = 3/4$ so that $v^\star = \ln 4 > v_0$.

of the payoff matrix decrease and as a result, the expected downstream payoff, $\tilde{\pi}_i(V_i)$, and welfare decrease. On the other hand, revenue increases in a more competitive environment as the data seller can threaten each data buyer to provide exclusive information to their rivals.

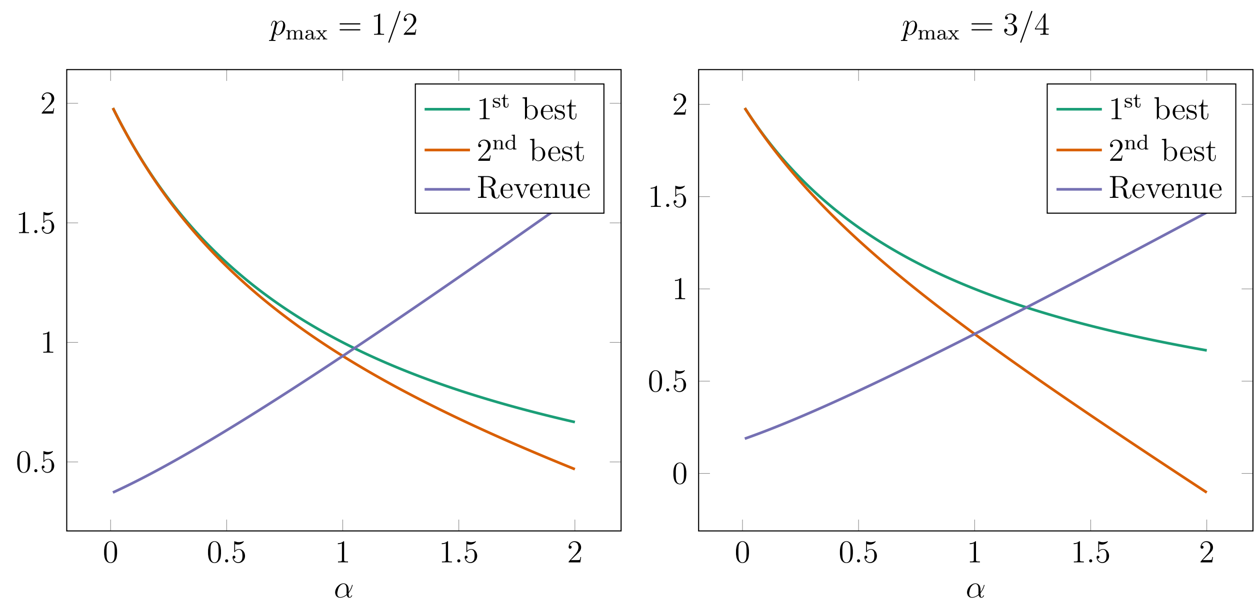

Figure 9: Welfare in the first best and second best mechanisms and revenue in the optimal mechanism as a function of α, for two different priors on the unknown state.

5.4 Number of Buyers

Another critical factor affecting the competition that a buyer faces in the downstream game is the total number of buyers. We now explore the impact of n on the optimal mechanisms.

In contrast to previous sections where we focused on the case $n = 2$ to visualize the impact of other factors, this requires considering the full generality of an n-dimensional type space, which is inherently hard to visualize. For this reason, we adopt the perspective of a single buyer $i \in [n]$ and study a two-dimensional slice of the type space parametrized by buyer i's type v_i, and the sum $s_{-i} = \sum_{j \neq i} v_j$ of the other buyers' types. Conveniently, the recommendation to buyer i in the welfare-optimal mechanism (Proposition 6) remains deterministic with this parametrization:

$$A_i = \theta \quad \text{if and only if} \quad s_{-i} \leq \max\{v^\star, (n-1)v_i/\alpha\}.$$

Similarly, buyer i's recommendation in the revenue-optimal mechanism (Proposition 9) is deterministic in $\phi(v_i)$ and $s^{\phi}_{-i} \coloneqq \sum_{j \neq i} \phi(v_j)$. In contrast, the externality induced by the recommendation to a buyer $j \neq i$ on buyer i remains random even after conditioning on v_i and s_{-i} and depends on the conditional distribution of v_j given s_{-i}.

Focusing on exponentially distributed type, s_{-i} follows an Erlang distribution, for which the conditional distribution $v_j \,|\, s_{-i}$ can be computed in closed-form. This allows us to visualize the externality induced by buyer j on buyer i by plotting the quantity $\mathbb{P}[A_j = \theta \mid V_i, S_{-i}]$ as a

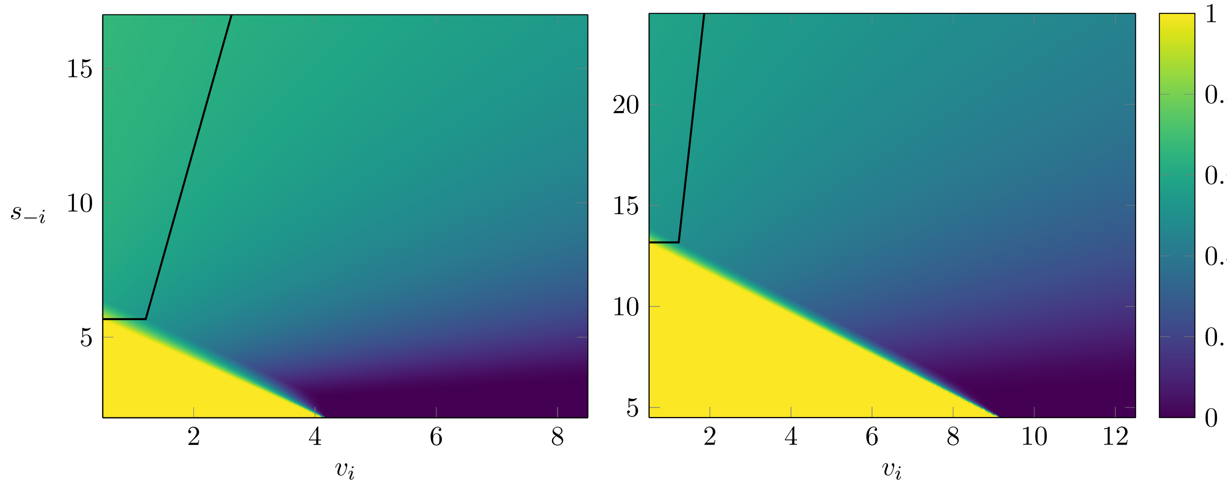

Figure 10: Heatmap of the function $(V_i, S_{-i}) \mapsto \mathbb{P}[A_j = \theta \mid V_i, S_{-i}]$ in the revenue-optimal mechanism for $n = 5$ (left) and $n = 10$ (right). Types are distributed as $1/2 + Y$ where Y is a standard exponential variable and $\alpha = P_{\max} = 1/2$. The solid black line shows the boundary determining the recommendation to buyer i: $A_i = \theta$ below the line and $A_i = 1 - \theta$ above it.

function of the two parameters (V_i, S_{-i}). Figure 10 shows a heatmap of this function in the revenue-optimal mechanism for two different values of n.

Finally, we turn to the impact of the number n of buyers on revenue and welfare. Specifically, we consider these objectives after normalization by n: by symmetry these correspond respectively to the utility (gross of any payments to the seller) and payment of a single buyer. As a way to assess the efficiency loss induced by the seller, we also consider the utility of a single buyer in the revenue-optimal mechanism. Figure 11 shows how these quantities vary with n for two different values of α. As we see, all three quantities converge to a constant as n grows to infinity. This can be heuristically explained as follows: because the externality term is normalized by $n - 1$ in (1), the optimal recommendation to buyer i, absent obedience constraints, is obtained by comparing their type v_i to the *average* of the other buyers' types. The law of large numbers then implies that in the large n limit, the situation that buyer i faces is identical to a competition with a *single* buyer whose type is concentrated on the mean of the type distribution. Consistent with our findings from the previous section, we also see that as α increases—competition becomes fiercer—the payments increase while buyers' utilities decrease.

6 Extensions

In this section, we present some extensions and variants of our model (Section 2) and show how to adapt our results to these. First, we explore the case of positive externalities in

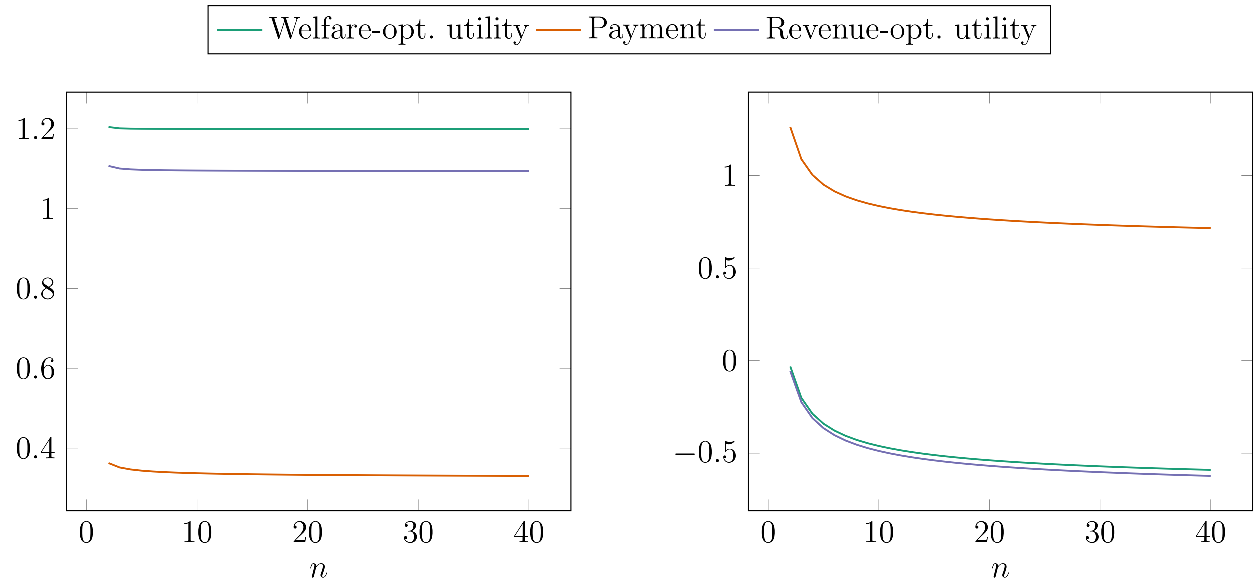

Figure 11: Utility of a buyer in the welfare- and revenue-optimal mechanisms (green and purple) and payment in the revenue-optimal mechanism (orange) for $\alpha = 0.2$ (left), and $\alpha = 2$ (right). Types are distributed as $1/2 + Y$ where Y is a standard exponential and $P_{\max} = 1/2$.

Section 6.1. Next, we consider two variants of the information structure: one in which the type of a buyer can be dependent on the state (Section 6.2) and one in which the buyers have the option to observe a private signal about the state after entering the mechanism (Section 6.3). Finally, we study the case where the seller is vertically integrated with one of the buyer in Section 6.4.

6.1 Strategic Complementarities

Although the class of games studied in the paper does not present complementarities strictly speaking, making the externality parameter α negative (resulting in positive externalities) could be interpreted as a form of strategic complementarities. Indeed in this case, if player i succeeds in matching their action to the state, they also benefit from player j matching the state and hence their own action. We have found that the optimal mechanisms can also be obtained with $\alpha < 0$ with only minor changes to our current analysis. Propositions 1 to 3, which pertain to characterizations of incentive compatibility, truthfulness, and obedience respectively, as well as Proposition 5 remain unchanged. Hence, the only necessary adjustments are in Proposition 6 (for welfare) and Proposition 9 (for revenue), which require dividing by α (now negative) in the inequality resulting from Proposition 5. To ensure the implementability of the mechanism, we need to verify that the induced interim payoff $\tilde{\pi}_i$ is non-decreasing.

Proposition 11. *Consider the binary game with additive payoffs (1), under Assumption 1* *and with $\alpha < 0$. The welfare-optimal mechanism sends the correct action recommendation to every data buyer regardless of their value. This mechanism is also implementable.*

Proof. By Proposition 5, the welfare-optimal recommendation subject to obedience is given by

$$A_i = \theta \quad \text{if and only if} \quad w_i(v) \geq \min\{\,0, t_i^\star(v_i)\,\},$$

where here $w_i(v) = v_i - \overline{\alpha} \sum_{j \neq i} v_j$. Since $\overline{\alpha}$ is negative and the types are non-negative, we have $w_i(v) \geq 0$ for all v and the optimal recommendation simplifies to $A_i = \theta$ for all v. The interim payoff of player i is constant and equal to $1 - \alpha$. Therefore, the the mechanism is implementable. □

Similarly, we can find the revenue-optimal mechanism.

Proposition 12. *Under the same assumptions as Proposition 11 and with the notations of Proposition 9, define $r^\star$ such that $F_\phi^{(n-1)}\big(\phi(r^\star)\big) = 1 - P_{\max}$. Then, the recommendation rule maximizing revenue subject to obedience is the deterministic rule given by*

$$A_i = \theta \quad \textit{if and only if} \quad \sum_{j \neq i} \phi(V_j) \geq \max\left\{\phi(r^\star), \frac{\phi(v_i)}{\overline{\alpha}}\right\}.$$

This recommendation rule is implementable as well.

Proof. By Lemma 8, finding the revenue-optimal mechanism reduces to maximizing virtual surplus where we have $w_i(v) = \phi(v_i) - \overline{\alpha} \sum_{j \neq i} \phi(v_j)$ as the weight functions. By Proposition 5, we compute the threshold function $t_i^\star(v_i)$ by solving

$$\mathbb{P}[w_i(v_i, V_{-i}) \geq t_i^\star(v_i)] = P_{\max}.$$

We obtain $t_i^\star(v_i) = \phi(v_i) - \overline{\alpha}\phi(r^\star)$. The allocation rule for the optimal revenue mechanism then becomes

$$A_i = \theta \quad \text{if and only if} \quad \sum_{j \neq i} \phi(V_j) \geq \max\left\{\phi(r^\star), \frac{\phi(v_i)}{\overline{\alpha}}\right\}.$$

Similarly, the expression for the interim payoff $\tilde{\pi}_i(v_i)$ can be derived as:

$$\tilde{\pi}_i(v_i) = 1 - F_\phi^{(n-1)}\left(\max\left\{\phi(r^\star), \frac{\phi(v_i)}{\overline{\alpha}}\right\}\right) - \overline{\alpha}\left(1 - F_\phi^{(n-2)}\left(\max\left\{\phi(r^\star), \frac{\phi(v_i)}{\overline{\alpha}}\right\} - \phi(v_i)\right)\right),$$

which is non-decreasing and therefore implementable. □

6.2 State-dependent Payoff Types

We explore the case where the buyers' payoff types are allowed to depend on the state. For example if θ represents the consumers' preference, the profitability of matching it could be larger in one state compared to the other. We thus generalize Assumption 1 by writing the utility of buyer $i \in [n]$ as

$$u_i(a; \theta, v_i) = v_i \cdot \beta_\theta \cdot \pi_i(a; \theta) \tag{10}$$

where β_θ is a commonly known, state-dependent (positive) scaling of the buyer's type.

The following proposition extends our characterization of obedience to utilities of the form (10). We still assume a binary state and separable externalities.[13]

Proposition 13. *For the binary game with additive payoffs* (1) *and assuming* (10)*, a recommendation rule is obedient if and only if the following inequality holds almost surely for each buyer* $i \in [n]$,

$$\mathbb{E}[\beta_\theta \mathbf{1}\{\theta = A_i\} \mid V_i] \geq \max_{k \in \{0,1\}} \beta_k \mathbb{P}[\theta = k].$$

Proof. We follow the same steps as in the proof of Proposition 3. Because externalities are separable, they cancel out when comparing following the recommendation to deviating from it. Obedience of player $i \in [n]$ is thus equivalent to the following two inequalities

$$\beta_1 \mathbb{P}[\theta = 1 \wedge A_i = 1 \mid V_i] \geq \beta_0 \mathbb{P}[\theta = 0 \wedge A_i = 1 \mid V_i]$$
$$\beta_0 \mathbb{P}[\theta = 0 \wedge A_i = 0 \mid V_i] \geq \beta_1 \mathbb{P}[\theta = 1 \wedge A_i = 0 \mid V_i].$$

Adding the left-hand side of the second (resp. first) inequality to the first (resp. second), we obtain

$$\beta_0 \mathbb{P}[\theta = 0 \wedge A_i = 0 \mid V_i] + \beta_1 \mathbb{P}[\theta = 1 \wedge A_i = 1 \mid V_i] \geq \beta_0 \mathbb{P}[\theta = 0]$$
$$\beta_0 \mathbb{P}[\theta = 0 \wedge A_i = 0 \mid V_i] + \beta_1 \mathbb{P}[\theta = 1 \wedge A_i = 1 \mid V_i] \geq \beta_1 \mathbb{P}[\theta = 1],$$

where we recognize the same quantity on the left-hand side of both inequalities. Hence, these two inequalities are equivalent to

$$\beta_0 \mathbb{P}[\theta = 0 \wedge A_i = 0 \mid V_i] + \beta_1 \mathbb{P}[\theta = 1 \wedge A_i = 1 \mid V_i] \geq \max_{k \in \{0,1\}} \beta_k \mathbb{P}[\theta = k]. \qquad \square$$

The characterization in Proposition 13 is similar to the one obtained in the paper. The quantity on the right-hand side quantifies the downstream payoff (ignoring externalities) that a buyer can guarantee themself by playing optimally with respect to the prior distribution.

[13] This result also holds for more general externality models, cf. Footnote 7.

Hence, this characterization still expresses that the recommendation rule needs to be "correct enough" to guarantee the buyer a utility at least as good as what they could get under the prior. The main difference being that the correctness of the recommendation $\mathbf{1}\{\theta = A_i\}$ is now scaled by the multiplier β_θ. Of course, when $\beta_0 = \beta_1 = 1$, we have $\mathbb{E}[\beta_\theta \mathbf{1}\{\theta = A_i\} \mid V_i] = \mathbb{P}[\theta = A_i \mid V_i]$ and we recover the characterization of Proposition 3 in the paper.

A *measure change* perspective. We can obtain an alternative proof of Proposition 13, as well as a convenient way to derive its implications for the design of optimal mechanisms as follows. We normalize β_θ so that its expectation is 1. That is, we define

$$\tilde{\beta}_\theta \coloneqq \frac{\beta_\theta}{\mathbb{E}[\beta_\theta]} = \frac{\beta_\theta}{\sum_{k\in\{0,1\}} \beta_k \mathbb{P}[\theta = k]} .$$

Note that we could replace β_θ with $\tilde{\beta}_\theta$ in (10) since this amounts to rescaling buyers' utilities by a constant. Now, $\tilde{\beta}_\theta$ can be interpreted as the probability density of a probability measure $\tilde{\mathbb{P}}$ over the state θ with respect to the original prior $\mathbb{P}$. In other words, $\tilde{\mathbb{P}}$ is defined by

$$\tilde{\mathbb{P}}[\theta = k] = \tilde{\beta}_k \mathbb{P}[\theta = k].$$

If we denote by $\tilde{\mathbb{E}}$ the expectations computed according to the new prior $\tilde{\mathbb{P}}$, we have in particular

$$\mathbb{E}[u_i(a;\theta, v_i) = \mathbb{E}[v_i \cdot \tilde{\beta}_\theta \cdot \pi_i(a;\theta)] = \tilde{\mathbb{E}}[v_i \cdot \pi_i(a;\theta)].$$

In other words, *the generalized model Equation* (10) *in which utilities are scaled with β_θ reduces to the model studied in our paper with the prior $\mathbb{P}$ replaced with $\tilde{\mathbb{P}}$.*

From this, we immediately obtain as a characterization of obedience

$$\tilde{\mathbb{P}}[\theta = A_i \mid V_i] \geq \tilde{P}_{\max} \coloneqq \max_{k\in\{0,1\}} \tilde{\mathbb{P}}[\theta = k],$$

which is equivalent to the statement in Proposition 13 after replacing $\tilde{\mathbb{P}}$ with its definition.

Similarly, the welfare-optimal mechanism keeps the exact same form as the one given in Proposition 6, with the only difference being that the critical type $v^\star$ is now defined by

$$F^{(n-1)}(v^\star) = \tilde{P}_{\max},$$

and similarly for the revenue-optimal mechanism.[14]

[14] For a lower level argument that does not require a change of measure, the reader can verify that all the steps in Section 4.1 in the paper still go through with the new characterization in Proposition 13, after defining $h_i(V) = \mathbb{E}[\beta_\theta \mathbf{1}\{\theta = A_i\} \mid V]$ instead of $h_i(V) = \mathbb{P}[\theta = A_i \mid V]$.

Remark 4. As we see, the more profitable a state is (larger β_θ) the more skewed the resulting β-tilted prior $\tilde{\mathbb{P}}$ is. Intuitively, if a state is so profitable as to outweigh the uncertainty, playing the action matching this state is a good strategy for the buyer. Consequently, a recommendation rule needs to reveal the correct action more often in order to be obedient. This extends our discussion of the buyer's prior information in Section 5.2.

6.3 Partially Informed Buyers

In many situations, the buyers may be able to access side information about the state θ. In our motivating example of firms acquiring information about consumer preferences in a downstream market, this could for example be the case if each firm has a marketing department that conducts its own market study in addition to the information acquired from the seller. Such situations can be coarsely classified into three categories:

1. Each buyer observes a signal about the state before entering the mechanism, and the buyers' signals are known to the seller. Upon receiving their signal, each buyer performs a Bayesian update of the common prior over θ. In our setting, this amounts to saying each buyer assigns a different (but known to the seller) probability to the most likely state. In other words the quantity $P_{\max}$ appearing our characterization of obedience (Proposition 3) is now buyer-dependent. Our results can be easily adapted to such situations at the cost of symmetry in the notations: the threshold type $v^\star$ guaranteeing obedience in our optimal mechanisms (Propositions 6 and 9) simply needs to be defined with respect to each buyer's individual $P_{\max}$. This setting was also studied by Bergemann and Morris (2013, Sec. 3.3) for continuous state and action spaces.

2. Each buyer observes a private signal about the state before entering the mechanism, but these signals are unknown to the seller who needs to elicit them. This setting was studied in matching game similar to ours in the case of a single buyer by Bergemann et al. (2018). With more than one buyer, this is an especially challenging problem due to the difficulty of eliciting correlated types. This setting is studied by Rodríguez Olivera (2021) in the case of a binary type space.

3. The buyers observe private signals about the state *after* entering the mechanism (but before taking actions in the downstream game). Prior to entering the mechanism, each buyer is privately informed about the *quality* of the signal they will receive at the second stage. The seller can thus attempt to elicit this information and condition the mechanism on it.

We focus on the third of the above categories and study a variant of our model in such a situation. To isolate the effect of the buyers' private information about the state from their intrinsic preference about the downstream payoffs, we make the simplifying assumption that their payoff types $(v_i)_{i\in[n]}$ are all identical and equal to 1. Consequently, the utility of buyer $i \in [n]$ is simply

$$u_i(a;\theta) = \pi_i(a;\theta),$$

where π_i is the downstream payoff of buyer i as introduced in our paper. The private type of buyer i—the quality of the private signal they will observe in the second stage of the game—will be denoted by t_i. We assume that $t_i \in [0,1]$ and interpret it as the buyer's probability of correctly guessing the correct action *after* observing their private signal at the second stage of the game. In other words, if A_i denotes the action recommendation from the mechanism, the buyer only values it if $\mathbb{P}[A_i = \theta \mid t_i] > t_i$. Otherwise, the recommendation is *uniformative* from the perspective of the buyer who will (weakly) prefer to ignore the mechanism's recommendation and play at the second stage according to their private signal. We observe that crucially, the buyers' types do not directly enter their utility, but only affect how they will react to the mechanism's recommendations.

By the revelation principle, we focus on truthful mechanisms in which each buyer makes a type report $t_i' \in [0,1]$. Note however than in contrast to the model studied in our paper, we cannot restrict ourselves to obedient mechanisms without loss of generality. This is because the mechanism does not have access to the agents' private signals and hence cannot send an action recommendation that "simulates" what the agent would have done after observing their signal. The following proposition characterizes such incentive compatible mechanisms. Recall that $\tilde{\pi}_i$ denotes buyer i's interim expected utility: $\tilde{\pi}_i(t_i) = \mathbb{E}[\pi_i(A;\theta) \mid t_i]$ when truthfully reporting their type t_i.

Proposition 14. *A mechanism is incentive compatible iff for each agent i*

1. *there exists a threshold type $t_i^\star$ such that the mechanism's recommendation to buyer i is uninformative for types $t_i > t_i^\star$.*

2. *there exists a constant C_i such that the payment function p_i is given by*

$$p_i(t_i) = \tilde{\pi}_i(t_i) - \max\{t_i, t_i^\star\} + C_i.$$

Proof. Let us write $\tilde{\pi}_i(t_i'; t_i)$ the expected interim utility of buyer i when their true type is t_i

but they report type t_i'. We have[15]

$$\tilde{\pi}_i(t_i'; t_i) = \max\{t_i, \mathbb{P}[A_i = \theta \mid t_i']\} + \tilde{E}_i(t_i'),$$

where we write $\tilde{E}_i(t_i')$ for the interim expected externality incurred by buyer i when reporting t_i'.

Let t_i' and t_i be two types for which the mechanisms' recommendation is informative, then truthfulness implies

$$\begin{aligned}
\tilde{\pi}_i(t_i; t_i) - p_i(t_i) &\geq \tilde{\pi}_i(t_i'; t_i) - p_i(t_i') = \max\{t_i, \mathbb{P}[A_i = \theta \mid t_i']\} + \tilde{E}_i(t_i') - p_i(t_i') \\
&\geq \mathbb{P}[A_i = \theta \mid t_i'] + \tilde{E}_i(t_i') - p_i(t_i') \\
&\geq \tilde{\pi}_i(t_i'; t_i') - p_i(t_i').
\end{aligned}$$

After swapping the role of t_i' and t_i, we get that $t_i \mapsto \tilde{\pi}_i(t_i) - p_i(t_i)$ must be constant over the set of all types that receive an informative recommendation. Let us denote this constant by D_i.

Similarly, by considering two types that receive an uninformative recommendation, we get that $t_i \mapsto \tilde{E}_i(t_i) - p_i(t_i) = \tilde{\pi}_i(t_i) - t_i - p_i(t_i)$ is constant over all such types. Let $\bar{D}_i$ denote this constant.

If t_i is an informed type and t_i' is an uniformed type, truthfulness implies

$$D_i \geq t_i + \bar{D}_i,$$

and similarly if t_i is uninformed while t_i' is informed

$$t_i + \bar{D}_i \geq D_i.$$

From which we immediately obtain the existence of the claimed threshold type $t_i^\star$ with $t_i^\star = D_i - \bar{D}_i$ and the stated expression for the payment with $C_i = -\bar{D}_i$ □

Proposition 15. *Define for $t \in [0,1]$, $G(t) := \int_t^1 s \mathrm{d}F(s) = 1 - tF(t) - \int_t^1 F(s)\mathrm{d}s$, and let $t^\star \in [0,1]$ be such that*

$$t^\star \in \operatorname*{arg\,max}_{t \in [0,1]} (1-t)F(t) + \alpha\big(1 - F(t) - G(t)\big).$$

[15] Note that this expression implicitly assumes that the agent commits to only observing one signal: either the seller's action recommendation or their private signal, but not both. Otherwise, we would also need to take into account the signals' correlation and the fact that the agent could be correct more often by combining the two signals than which each signal separately. This seems to lead to an untractable problem.

When $\alpha < 1$, the mechanism that maximizes revenue subject to truthfulness recommends the correct action ($A_i = \theta$) to all types $t_i \leq t_i^\star$ and gives an uninformative recommendation otherwise. The payment collected from player i is

$$p_i(t_i) = (1 - t^\star)\mathbf{1}\{t_i \leq t^\star\} + \alpha\big(1 - F(t^\star) - G(t^\star)\big).$$

Proof. It follows immediately from Proposition 14 that the expected revenue R of any incentive compatible mechanism is given by

$$R = W - \sum_{i \in [n]} \big(t_i^\star F(t_i^\star) + G(t_i^\star) - C_i\big)$$

where W is the expected welfare. Furthermore, we easily get that the optimal choice of C_i that still guarantees participation is $C_i = \alpha$.

Thus for a fixed threshold type $t_i^\star$, the problem reduces to maximizing welfare (subject to sending uninformative recommendations for types $t_i \geq t_i^\star$). So we solve this optimization problem for fixed $t_i^\star$ and optimize over $t_i^\star$ at the end. Since,

$$W = (1 - \alpha) \sum_{i \in [n]} \mathbb{P}[A_i = \theta]$$

we see that when $\alpha < 1$ we want to maximize the probability that the recommendation is correct, which we achieve by always recommending the correction action when $t_i \leq t_i^\star$. For this mechanism the expected payment collected from buyer i is

$$p_i(t_i) = (1 - t_i^\star)F(t_i^\star) + \alpha\big(1 - F(t_i^\star) - G(t_i^\star)\big).$$

The revenue-optimal is thus obtained by finding the threshold $t_i^\star$ that maximizes the previous expression, as stated in the proposition statement. □

An equivalent way to describe the revenue-optimal mechanism is that it is a combination of a fixed fee and a posted price mechanism. The mechanism always "extorts" the fixed fee $\alpha(1 - F(t^\star) - G(t^\star))$, which is exactly the amount of averted externality by having everyone participate. In addition, the seller offers to reveal the correct state θ at the price $1 - t_i^\star$. The buyer only accepts the offer if the gain $(1 - t_i)$ exceeds the cost, or equivalently if $t_i \leq t_i^\star$.

6.4 Vertical Integration

We now study two situations in which the information seller vertically integrates with one of the buyers and describe in both how this modifies the optimal mechanisms. We focus on the case $n = 2$ for simplicity, but the general case does not present any additional difficulty.

(A) In the first scenario, the information seller is one of the two firms, say, firm 1, both observing the state θ *and* participating in the downstream game against firm 2 (the data buyer). The information seller designs a payment function $p_2 : \mathcal{V}_2 \to \mathbb{R}_{\geq 0}$ and communication rule $\sigma : \Theta \times \mathcal{V} \to \Delta(\mathcal{A}_2)$. Firm 2 pays $p_2(v_2)$ to firm 1 in exchange for an action recommendation a_2 distributed according to $\sigma(\theta, v)$, resulting in net utility $v_2\pi_2(a;\theta) - p_2(v_2)$. Firm 1's utility is $v_1\pi_1(a;\theta) + p_2(v_2)$, and since it knows the state θ, it will play the dominant strategy $a_1 = \theta$. In this scenario, the incentive compatibility and individual rationality constraints of firm 2 remain identical to Section 3.1. In particular, one can derive the payment p_2 from the interim payoff $\tilde{\pi}_2$ according to Proposition 2.

Maximizing welfare reduces to maximizing $\mathbb{E}\big[(V_2 - \alpha V_1)\mathbf{1}\{A_2 = \theta\}\big]$, which we achieve by following the same steps as in Proposition 6. The optimal recommendation is given by

$$A_2 = \theta \quad \text{if and only if} \quad v_1 \leq \max\{v^\star, v_2/\alpha\},$$

where $v^\star$ is defined by $F(v^\star) = \mathbb{P}[V_1 \leq v^\star] = P_{\max}$. This recommendation is identical to the one firm 2 would have received in our original model without integration (Proposition 6). The only difference from the perspective of firm 2 is that firm 1 now plays the dominant action with probability one, inducing as interim downstream payoff

$$\tilde{\pi}_2(v_2) = \max\{P_{\max}, F(v_2/\alpha)\} - \alpha, \tag{11}$$

in which the externality $-\alpha$ is worse than in Proposition 7. Since $\tilde{\pi}_2$ is non-decreasing, the mechanism is implementable.

If firm 1 maximizes its own utility instead of welfare, the optimization problem reduces to maximizing $\mathbb{E}\big[\big(\phi(V_2) - \alpha V_1\big)\mathbf{1}\{A_2 = \theta\}\big]$. Assuming that V_2's distribution is regular, the virtual value function ϕ is non-decreasing and the optimal recommendation is given by

$$A_2 = \theta \quad \text{if and only if} \quad v_1 \leq \max\{v^\star, \phi(v_2)/\alpha\}$$

where $v^\star$ is defined as above. The associated interim payoff is

$$\tilde{\pi}_2(v_2) = \max\{P_{\max}, F\big(\phi(v_2)/\alpha\big)\} - \alpha,$$

which is non-decreasing in v_2 and thus implementable.

(B) In the second scenario, the information seller and firm 1 remain distinct entities: in particular firm 1 still needs to acquire information and the seller does not engage in the downstream game. However, the seller and firm 1 are vertically integrated in the sense that the former wishes to maximize the sum of firm 1's utility and his own revenue.

In this scenario, the seller's constraints—truthfulness, obedience, and individual rationality—stay the same as in Section 3.1. The seller's objective function becomes

$$\mathbb{E}\big[V_1\pi_1(A;\theta) - p_1(\theta, V) + p_1(\theta, V) + p_2(\theta, V)\big] = \mathbb{E}\big[V_1\pi_1(A;\theta) + p_2(\theta, V)\big]$$

and by Lemma 8, maximizing it reduces to maximizing

$$\mathbb{E}\Big[\big(V_1 - \alpha\phi(V_2)\big)\mathbf{1}\{A_1 = \theta\} + \big(\phi(V_2) - \alpha V_1\big)\mathbf{1}\{A_2 = \theta\}\Big].$$

The optimal recommendation rule is obtained using Proposition 5 and we find

$$\begin{aligned} A_1 = \theta \quad &\text{if and only if} \quad \phi(v_2) \le \max\{\phi(v^\star), v_1/\alpha\} \\ A_2 = \theta \quad &\text{if and only if} \quad v_1 \le \max\{v^\star, \phi(v_2)/\alpha\}, \end{aligned} \tag{12}$$

where $F(v^\star) = P_{\max}$. Compared to the revenue-optimal recommendation rule of Proposition 9 where we had

$$A_i = \theta \quad \text{if and only if} \quad \phi(v_j) \le \max\{\phi(v^\star), \phi(v_i)/\alpha\},$$

we see that buyer 1 receives the correction recommendation more often (since $\phi(x) \le x$). In contrast, buyer 2 receives the correct recommendation less often in (12) compared to Proposition 9, due to $\phi^{-1}(x) \ge x$. This is expected given that the seller's objective now favors firm 1.

We can also compute the interim payoffs from (12) and find

$$\begin{aligned} \tilde{\pi}_1(v_1) &= \max\{P_{\max}, F \circ \phi^{-1}(v_1/\alpha)\} - \alpha\Big[1 - \mathbf{1}\{v_1 > v^\star\}F \circ \phi^{-1}(\alpha v_i)\Big] \\ \tilde{\pi}_2(v_2) &= \max\{P_{\max}, F\big(\phi(v_2)/\alpha\big)\} - \alpha\Big[1 - \mathbf{1}\{v_2 > v^\star\}F\big(\alpha\phi(v_2)\big)\Big], \end{aligned}$$

which are both non-decreasing since F and ϕ are both non-decreasing. The recommendation rule (12) is therefore implementable.

7 Conclusions

We have explored the implications of selling information to competing buyers in a mechanism design framework. The nature of information disciplines the optimal mechanisms for selling data products and distinguishes them from canonical (e.g., physical) goods. In particular, the buyers' actions in the downstream game introduce obedience constraints into the designer's choice of mechanism. These constraints prevent a social planner from implementing the efficient degree of information *exclusivity*: the second-best mechanism involves symmetric allocations of information more often than optimal. At the same time, obedience also limits the allocation distortions introduced by a monopolist seller of information: the revenue-optimal mechanism provides the correct information to the buyers more often than the monopolist would like.

In the present work, we characterized optimal mechanisms in the context of a linear model with binary states and actions. Considerable work remains to be done to extend the applicability of this framework to real-world markets. Natural next steps include introducing asymmetric buyers and externality parameters, which could be represented, for example, by a weighted directed graph. Both extensions can be analyzed in our linear model. Several other extensions can be accommodated in our model but are unlikely to yield a tractable analysis. These include the possibility of downstream resale of information, the presence of competing information sellers, and the sequential arrival of information buyers.

Finally, in many models of downstream competition, information generates *nonlinear* externalities that also depend on all buyers' actions in the downstream game. In these settings, the sale of information to competing buyers creates value not only by allowing buyers to match their actions to the state but also by enabling coordination. In ongoing work (Bonatti et al., 2022), we pursue the coordination role of selling information in a linear-quadratic *Gaussian* model.

A Appendix

Proof of Proposition 1. Consider a truthful and obedient mechanism. Then, using the notations of Definition 1, we have for each $(v_i, v_i') \in \mathcal{V}_i$:

$$
\begin{aligned}
\mathbb{E}\big[u_i(A;\theta,V) - p_i(v_i)|V_i = v_i, B_i = v_i] \\
&= \mathbb{E}\big[v_i \cdot \pi_i(A;\theta) - p_i(v_i)|V_i = v_i, B_i = v_i] \\
&\geq \mathbb{E}\big[v_i \cdot \pi_i(A;\theta) - p_i(v_i')|V_i = v_i, B_i = v_i'] \\
&= \mathbb{E}\big[v_i \cdot \pi_i(A;\theta) - p_i(v_i')|B_i = v_i'] \\
&= v_i \cdot \mathbb{E}\big[\pi_i(A;\theta) \mid B_i = v_i'] - \mathbb{E}[p_i(v_i')|B_i = v_i'],
\end{aligned}
\tag{13}
$$

where the first equality is by Assumption 1, the inequality is by truthfulness (Definition 3), and the third equality is because A is independent of V_i conditioned on B_i by Assumption 1.

Let $\delta : A_i \to A_i$ be a deviation function. By Assumption 1, obedience (Definition 2) is equivalent to

$$
\mathbb{E}[\pi_i(A;\theta) \mid B_i = v_i] \geq \mathbb{E}[\pi_i(\delta(A_i), A_{-i};\theta) \mid B_i = v_i]
$$

for all $v_i \in \mathcal{V}_i$. Applying this inequality for v_i' in (13) we obtain

$$
\begin{aligned}
\mathbb{E}\big[u_i(A;\theta,V) - p_i(v_i)|V_i = v_i, B_i = v_i] &\geq v_i \cdot \mathbb{E}\big[\pi_i(\delta(A_i), A_{-i},\theta) \mid B_i = v_i'] \\
&\quad - \mathbb{E}[p_i(v_i')|B_i = v_i'] \\
&\geq \mathbb{E}\big[u_i(\delta(A_i), A_{-i};\theta,V) \mid V_i = v_i, B_i = v_i'] \\
&\quad - \mathbb{E}[p_i(v_i')|V_i = v_i, B_i = v_i'],
\end{aligned}
$$

which is precisely the definition of incentive compatibility. □

Proof of Proposition 2. Define $\tilde{u}_i : v_i \mapsto v_i \cdot \tilde{\pi}_i(v_i) - p_i(v_i)$. Then truthfulness is equivalent to

$$
\tilde{u}_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - p_i(v_i) \geq v_i \cdot \tilde{\pi}_i(v_i') - p_i(v_i') = \tilde{u}_i(v_i') + (v_i - v_i') \cdot \tilde{\pi}_i(v_i')\,,
$$

for all $(v_i, v_i') \in \mathcal{V}_i^2$. This is equivalent to saying that $\tilde{\pi}_i(v_i) \in \partial \tilde{u}_i(v_i)$ for all $v_i \in \mathcal{V}_i$ where $\partial \tilde{u}_i(v_i) \subset \mathbb{R}$ denotes the subdifferential of $\tilde{u}_i$ at v_i. By a well-known characterization of convexity, this in turn equivalent to saying that $\tilde{\pi}_i$ is non-decreasing and that

$$
\tilde{u}_i(v_i) = \tilde{u}_i(\underline{v}) + \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s)ds\,.
$$

This concludes the proof since this last expression is equivalent to (3). □

Proof of Proposition 3. Since there are two states, obedience states that for all $i \in [n]$ and all $a_i \in \Theta$,

$$\mathbb{E}[\pi_i(a_i, A_{-i}; \theta) - \pi_i(1 - a_i, A_{-i}; \theta) \mid A_i = a_i, V_i] \geq 0.$$

Using the form of π_i from Equation (1), we observe that the externality terms cancel out and the previous inequality is equivalent to

$$\mathbb{P}[\theta = a_i \mid A_i = a_i, V_i] \geq \mathbb{P}[\theta = 1 - a_i \mid A_i = a_i, V_i],$$

and by Bayes' rule to

$$\mathbb{P}[\theta = a_i \wedge A_i = a_i \mid V_i] \geq \mathbb{P}[\theta = 1 - a_i \wedge A_i = a_i \mid V_i]\,.$$

Adding the quantity $\mathbb{P}[\theta = 1 - a_i \wedge A_i = 1 - a_i \mid V_i]$ on both sides, we obtain

$$\mathbb{P}[A_i = \theta \mid V_i] \geq \mathbb{P}[\theta = 1 - a_i \mid V_i] = \mathbb{P}[\theta = 1 - a_i]\,,$$

where the last equality uses that θ and V_i are independent by Assumption 1. Taking a maximum over $a_i \in \Theta$ yields the lemma's statement. □

Proof of Lemma 4. Given functions h_i satisfying the lemma's assumptions, one can choose for example σ such that for all $(x, v) \in \Theta \times \mathcal{V}$, the distribution $\sigma(x, v) \in \Delta(\mathcal{A})$ has independent coordinates with marginals given by h_i. Formally, we have for $(x, v) \in \Theta \times \mathcal{V}$ and $a \in \mathcal{A}$,

$$\sigma(a; x, v) = \prod_{i : a_i = x} h_i(v) \prod_{i : a_i \neq x} \big(1 - h_i(v)\big).$$ □

Before we prove Proposition 5 we first state and prove a variational lemma that solves the pointwise optimization problem that the optimal mechanism reduces to.

Lemma 16 (Variational Lemma)**.** *Let (E, μ) be a probability space and let $g : E \to \mathbb{R}$ be a μ-integrable function whose level sets are μ-null sets: $\mu(\{v \in E \mid g(v) = k\}) = 0$ for all $k \in \mathbb{R}$. Consider the problem*

$$\begin{aligned} \max_{h \in \mathcal{F}} \ & \mathcal{L}(h) \coloneqq \int_E h \cdot g \, \mathrm{d}\mu \\ \textit{s.t.} \ & \int_E h \, \mathrm{d}\mu \geq c \end{aligned}$$

where the optimization is over the set $\mathcal{F}$ of measurable functions $h : E \to \mathbb{R}$ with $h(E) \subseteq [0, 1]$ and $c \in [0, 1]$ is a constant. For $k \in \mathbb{R}$, define $L_g^+(k) \coloneqq \{v \in E \mid g(v) \geq k\}$ the superlevel set of g of level k and let $t_c \coloneqq \sup\{k \in \mathbb{R} \mid \mu\big(L_g^+(k)\big) \geq c\}$. Then, an optimal solution to the *problem is given by* $h^\star : v \mapsto \mathbf{1}\{g(v) \geq t^\star\}$ *where* $t^\star = \min\{t_c, 0\}$.

Proof. First, note that t_c is well-defined in the extended real line by adopting the usual convention $\sup \emptyset = -\infty$ and $\sup \mathbb{R} = +\infty$. If $k \leq k'$, we have $L^+_{k'}(g) \subseteq L^+_k(g)$ showing that the function $m : k \mapsto \mu\big(L^+_k(g)\big)$ is non-increasing. The identity $L^+_g(k) = \bigcap_{k'<k} L^+_g(k')$ implies

$$m(k) = \mu\big(L^+_g(k)\big) = \inf_{k'<k} \mu\big(L^+_g(k')\big) = \inf_{k'<k} m(k')$$

and m is thus left-continuous. Consequently, the threshold t_c is characterized by the equivalence

$$k \leq t_c \iff \mu\big(L^+_g(k)\big) \geq c.$$

Furthermore, the identity $\{v \in E \mid g(v) > k\} = \bigcup_{k'>k} L^+_{k'}(g)$ implies

$$\mu\big(\{v \in E \mid g(v) > k\}\big) = \sup_{k<k'} \mu\big(L^+_g(k')\big) = \sup_{k<k'} m(k').$$

But because the level sets of g are μ-null sets, we also have $m(k) = \mu\big(\{v \in E \mid g(v) > k\}\big)$, hence the function m is continuous, and the threshold t_c further satisfies $\mu\big(L^+_g(t_c)\big) = c$. This implies that $h^\star$ is a feasible solution, indeed since $t^\star \leq t_c$ we have

$$\int_E h^\star \,\mathrm{d}\mu = \mu\big(L^+_g(t^\star)\big) \geq \mu\big(L^+_g(t_c)\big) = c.$$

Next, for a feasible $h \in \mathcal{F}$ we have

$$\begin{aligned}
\mathcal{L}(h^\star) - \mathcal{L}(h) &= \int_{L^+_g(t^\star)} (1-h)g \,\mathrm{d}\mu + \int_{E\setminus L^+_g(t^\star)} (-h)g \,\mathrm{d}\mu \\
&\geq t^\star \int_{L^+_g(t^\star)} (1-h) \,\mathrm{d}\mu + t^\star \int_{E\setminus L^+_g(t^\star)} (-h) \,\mathrm{d}\mu \\
&= t^\star \mu\big(L^+_g(t^\star)\big) - t^\star \int_E h \,\mathrm{d}\mu \geq t^\star \Big[\mu\big(L^+_g(t^\star)\big) - c\Big],
\end{aligned} \tag{14}$$

where the first equality is by definition of $h^\star$, the subsequent equality uses that h takes values in $[0,1]$ and the fact that $g(v) \geq t^\star$ iff $v \in L^+_g(t^\star)$, and the last inequality uses that $t^\star \leq 0$ by definition and that $\int_E h \,\mathrm{d}\mu \geq c$ by feasibility.

Either $t^\star = 0$, or $t^\star = t_c$ in which case $\mu\big(L^+_g(t^\star)\big) = c$. In both cases, the last expression in (14) vanishes, hence $\mathcal{L}(h^\star) \geq \mathcal{L}(h)$ for all feasible $h \in \mathcal{F}$ which concludes the proof. □

We can now describe the optimal mechanism.

Proof of Proposition 5. First note that since $w_i(v_i, V_{-i})$ is non-atomic, its c.d.f. is continuous for all $v_i \in \mathcal{V}_i$. Hence, defining $c \coloneqq \max_{k\in\{0,1\}} \mathbb{P}[\theta = k]$, a suitable threshold function $t_i^\star$ can be obtained by choosing for $v_i \in \mathcal{V}_i$,

$$t_i^\star(v_i) \coloneqq \sup\bigl\{k \in \mathbb{R} \mid \mathbb{P}[w_i(v_i, V_{-i}) \geq k] \geq c\bigr\}.$$

Define $h_i : \mathcal{V} \to [0,1]$ by $h_i(V) = \mathbb{P}[A_i = \theta \mid V]$. By Proposition 3 and using the law of iterated expectations, the obedience constraint for buyer $i \in [n]$ states that

$$\mathbb{P}[A_i = \theta \mid V_i] = \mathbb{E}[\mathbf{1}\{A_i = \theta\} \mid V_i] = \mathbb{E}\bigl[\mathbb{E}[\mathbf{1}\{A_i = \theta\} \mid V] \mid V_i\bigr] = \mathbb{E}[h_i(V) \mid V_i] \geq c$$

almost surely for V_i. Similarly we can rewrite the objective (5) in terms of the functions h_i and the optimization problem we need to solve can thus be written

$$\begin{aligned} \max \sum_{i=1}^{n} \mathbb{E}\bigl[w_i(V)h_i(V)\bigr] \\ \text{s.t. } \mathbb{E}[h_i(V) \mid V_i] \geq c, \text{ for } i \in [n]. \end{aligned}$$

Because both the objective function and the constraints are separable in i, this problem decomposes as n separate optimization problem, one for each h_i, $i \in [n]$:

$$\begin{aligned} \max \mathbb{E}\bigl[\mathbb{E}[w_i(V_i, V_{-i})h_i(V_i, V_{-i}) \mid V_i]\bigr] \\ \text{s.t. } \mathbb{E}[h_i(v_i, V_{-i})] \geq c, \text{ for all } v_i \in \mathcal{V}_i \end{aligned}$$

where we used the law of total expectation and the independence of $(V_1, \dots, V_n)$. Finally, since the constraint is a pointwise constraint for $v_i \in \mathcal{V}_i$, the optimal h_i is obtained by choosing the partial function $h_i(v_i, \cdot)$ so as to maximize the integrand in the objective function for each v_i. That is, $h_i(v_i, \cdot)$ should solve

$$\begin{aligned} \max \mathbb{E}\bigl[w_i(v_i, V_{-i})h_i(v_i, V_{-i})\bigr] \\ \text{s.t. } \mathbb{E}[h_i(v_i, V_{-i})] \geq c. \end{aligned}$$

This problem is exactly of the form solved in Lemma 16, with μ being the probability distribution of V_{-i} and with $g = w_i(v_i, \cdot)$ and $E = \mathbb{R}^{n-1}$. By construction, the threshold $t_i^\star(v_i)$ plays the role of t_c in Lemma 16 for this choice of function g. Hence, the optimal policy is given by

$$h_i(v_i, v_{-i}) = \mathbf{1}\bigl\{w_i(v_i, v_{-i}) \geq \min\{0, t_i^\star(v_i)\}\bigr\} \text{ for each } v_i \in \mathcal{V}_i \ ,$$

which is the deterministic rule in the proposition statement. $\square$

Then, the recommendation rule maximizing social welfare subject to obedience is the deterministic rule given by

$$A_i = \theta \quad \text{if and only if} \quad \sum_{j \neq i} v_j \leq \max\{v^\star, v_i/\overline{\alpha}\}.$$

Proof of Proposition 6. For the binary game with additive payoffs (1) under Assumption 1, the expected welfare (6) is of the form (5) covered by Proposition 5, with weight function $w_i(v) = v_i - \overline{\alpha} \sum_{j \neq i} v_j$. Furthermore, the level sets of the partial function $w_i(v_i, \cdot)$ are hyperplanes in dimension n and since the distribution of V_{-i} is absolutely continuous by assumption (with c.d.f. $F^{(n-1)}$), the random variable $w_i(v_i, V_{-i})$ is non-atomic for each $v_i \in \mathcal{V}_i$. Hence, Proposition 5 applies and the recommendation rule maximizing welfare subject to obedience is determined by

$$A_i = \theta \quad \text{if and only if} \quad w_i(v) \geq \min\{0, t_i^\star(v_i)\}.$$

All that remains is to compute the threshold function $t_i^\star(v_i)$. By definition we must have

$$\mathbb{P}\big[w_i(v_i, V_{-i}) \geq t_i^\star(v_i)\big] = \Pr\left[\sum_{j \neq i} V_j \leq \frac{v_i - t_i^\star(v_i)}{\overline{\alpha}}\right] = F^{(n-1)}\left(\frac{v_i - t_i^\star(v_i)}{\overline{\alpha}}\right) = P_{\max},$$

which by definition of $v^\star$ is equivalent to

$$\frac{v_i - t_i^\star(v_i)}{\overline{\alpha}} = v^\star.$$

The result in the statement follows after observing that $w_i(v) \geq \min\{0, t_i^\star(v_i)\}$ is equivalent to $\sum_{j \neq i} v_j \leq \max\{v^\star, v_i/\overline{\alpha}\}$. □

Proof of Proposition 7. Using the definition of the optimal recommendation rule and linearity of conditional expectations we have

$$\begin{aligned}
\tilde{\pi}_i(V_i) &= \mathbb{E}\Big[\mathbf{1}\{A_i = \theta\} - \overline{\alpha} \sum_{j \neq i} \mathbf{1}\{A_j = \theta\} \;\Big|\; V_i\Big] \\
&= \mathbb{E}\left[\mathbf{1}\left\{\sum_{j \neq i} V_j \leq \max\{v^\star, V_i/\overline{\alpha}\}\right\} \;\middle|\; V_i\right] - \overline{\alpha} \sum_{j \neq i} \mathbb{E}\left[\mathbf{1}\left\{\sum_{k \neq j} V_k \leq \max\{v^\star, V_j/\overline{\alpha}\}\right\} \;\middle|\; V_i\right].
\end{aligned}$$

The first expectation on the right-hand side can be computed as

$$\mathbb{E}\left[\mathbf{1}\left\{\sum_{j \neq i} V_j \leq \max\{v^\star, V_i/\overline{\alpha}\}\right\} \;\middle|\; V_i\right] = \max\big\{F^{(n-1)}(v^\star), F^{(n-1)}(V_i/\overline{\alpha})\big\}.$$

For $j \neq i$, the summand expectation can be computed using the law of total expectations as

$$\mathbb{E}\Big[\mathbf{1}\Big\{\sum_{k\neq j} V_k \leq \max\{v^\star, V_j/\overline{\alpha}\}\Big\} \;\Big|\; V_i\Big] = \mathbb{E}\Big[\mathbb{E}\Big[\mathbf{1}\Big\{\sum_{k\neq j} V_k \leq \max\{v^\star, V_j/\overline{\alpha}\}\Big\} \;\Big|\; V_i, V_j\Big] \;\Big|\; V_i\Big]$$
$$= \mathbb{E}\Big[F^{(n-2)}\big(\max\{v^\star, V_j/\overline{\alpha}\} - V_i\big) \;\Big|\; V_i\Big].$$

The claim that $\tilde{\pi}_i$ is a non-decreasing function easily follows from the observation that $v \mapsto \max\big\{F^{(n-1)}(v^\star), F^{(n-1)}(v/\overline{\alpha})\big\}$ is non-decreasing (because $F^{(n-1)}$ is non-decreasing as a c.d.f.) and that for all $j \neq i$, the quantity $\mathbb{E}\Big[F^{(n-2)}\big(\max\{v^\star, V_j/\overline{\alpha}\} - v\big)\Big]$, is non-increasing as a function of v because the integrand is non-increasing in v pointwise. □

Proof of Lemma 8.

1. Assume that p_i truthfully implements $\tilde{\pi}_i$, so that the interim utility $\tilde{u}_i$ is convex by Proposition 2. If (σ, p) is individually rational, then it is obviously individually rational at the lowest type. For the converse direction, assume that $\tilde{u}_i(\underline{v}) \geq \underline{v} \cdot K$, then we have for all $v_i \in \mathcal{V}_i$

   $$\tilde{u}_i(v_i) \geq \tilde{u}_i(\underline{v}) + (v_i - \underline{v}) \cdot \tilde{\pi}_i(\underline{v}) \geq v_i \cdot K\,,$$

   where the first inequality uses convexity of $\tilde{u}_i$ and the second inequality uses that $\tilde{\pi}_i(\underline{v}) \geq K$ and individual rationality at the lowest type.

2. If p_i implements $\tilde{\pi}_i$ truthfully, then we already know by Proposition 2 that it satisfies (3), in which the only undetermined quantity is the payment at the lowest type $p_i(\underline{v})$. By 1., individual rationality further constrains $p_i(\underline{v}) \leq \underline{v} \cdot (\tilde{\pi}_i(\underline{v}) - K)$, so the revenue-maximizing choice makes this constraint bind and we get (7).

   For this payment function, we write the seller's revenue as

   $$R = \mathbb{E}\Big[\sum_{i=1}^{n} p_i(V_i)\Big] = \sum_{i=1}^{n} \mathbb{E}\Big[V_i \cdot \tilde{\pi}_i(V_i) - \underline{v} \cdot K - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s) ds\Big]$$
   $$= \sum_i \mathbb{E}\Big[V_i \cdot \tilde{\pi}_i(V_i) - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s) ds\Big] - n\underline{v} \cdot K. \tag{15}$$

The expectation in the first summand is computed as follows

$$\begin{aligned}
\mathbb{E}\Big[V_i \cdot \tilde{\pi}_i(V_i) - \int_{\underline{v}}^{V_i} \tilde{\pi}_i(s)ds\Big] &= \int_{\underline{v}}^{\overline{v}} v_i \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i - \int_{\underline{v}}^{\overline{v}} \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) f(v_i) ds dv_i \\
&= \int_{\underline{v}}^{\overline{v}} v_i \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i - \int_{\underline{v}}^{\overline{v}} \int_{s}^{\overline{v}} \tilde{\pi}_i(s) f(v_i) dv_i ds \\
&= \int_{\underline{v}}^{\overline{v}} \left(v_i - \frac{1 - F(v_i)}{f(v_i)} \right) \cdot \tilde{\pi}_i(v_i) f(v_i) dv_i \\
&= \mathbb{E}\big[\phi(V_i)\tilde{\pi}_i(V_i)\big] ,
\end{aligned}$$

where the second equality uses Fubini's theorem and the last equality is the definition of the virtual value function. □

Proof of Proposition 10. The computation of the interim payoff $\tilde{\pi}_i$ is identical to the one in the proof of Proposition 7 after replacing the types with their virtual counterparts. We obtain for the interim downstream payoff

$$\begin{aligned}
\tilde{\pi}_i(v_i) = \max \left\{ F_{\phi}^{(n-1)} \left(\phi(v^\star)\right), F_{\phi}^{(n-1)} (\phi(v_i)/\overline{\alpha}) \right\} \\
- \overline{\alpha} \sum_{j \neq i} \mathbb{E}\left[F_{\phi}^{(n-2)} \left(\max \left\{ \phi(v^\star), \phi(V_j)/\overline{\alpha} \right\} - \phi(v_i) \right) \right] ,
\end{aligned}$$

which is non-decreasing assuming that F is regular (i.e., ϕ non-decreasing). This guarantees truthfulness of the mechanism.

It is immediate from this expression that for all $v_i \in \mathcal{V}_i$,

$$\tilde{\pi}_i(v_i) \geq F_{\phi}^{(n-1)}\big(\phi(v^\star)\big) - \alpha = P_{\max} - \alpha,$$

where the inequality is obtained by upper-bounding $F_{\phi}^{(n-2)}$ by 1 and the equality is by definition of $v^\star$. Furthermore, for the outside allocation described in the proposition statement, the best response of buyer i in case of non-participation is to play the action matching the most likely state under the prior, resulting in an expected payoff $K := P_{\max} - \alpha$.

Indeed, the non-participating buyer will be correct with probability $P_{\max}$ while incurring an externality of $-\alpha$ from all the participating buyers (who then receive the correct action recommendation). The previous two equations combined show that $\tilde{\pi}_i(v_i) \geq K$ and the conditions of Lemma 8 are thus satisfied. The revenue-maximizing payments subject to truthfulness and individual rationality are then given by Equation (7). □

References

Admati, A. R. and Pfleiderer, P. (1986). A monopolistic market for information. *Journal of Economic Theory*, 39(2):400–438.

Admati, A. R. and Pfleiderer, P. (1990). Direct and indirect sale of information. *Econometrica*, 58(4):901–928.

Agarwal, A., Dahleh, M., Horel, T., and Rui, M. (2020). Towards data auctions with externalities. *arXiv preprint arXiv:2003.08345*.

Babaioff, M., Kleinberg, R., and Paes Leme, R. (2012). Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 92–109.

Bergemann, D. and Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11:85–107.

Bergemann, D., Bonatti, A., and Smolin, A. (2018). The design and price of information. *American Economic Review*, 108(1):1–48.

Bergemann, D. and Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308.

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bergemann, D. and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Bimpikis, K., Crapis, D., and Tahbaz-Salehi, A. (2019). Information sale and competition. *Management Science*, 65(6):2646–2664.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272.

Bonatti, A., Dahleh, M., Horel, T., and Nouripour, A. (2022). Coordination via selling information. Technical report, MIT.

Bounie, D., Dubus, A., and Waelbroeck, P. (2021). Selling strategic information in digital competitive markets. *RAND Journal of Economics*, 52(2):283–313.

Jehiel, P. and Moldovanu, B. (2000). Auctions with downstream interaction among buyers. *RAND journal of economics*, 31(4):768–791.

Jehiel, P. and Moldovanu, B. (2006). Allocative and informational externalities in auctions and related mechanisms. In Blundell, R., Newey, W., and Persson, T., editors, *Proceedings of the 9th World Congress of the Econometric Society*. Cambridge University Press.

Jehiel, P., Moldovanu, B., and Stacchetti, E. (1996). How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–829.

Jehiel, P., Moldovanu, B., and Stacchetti, E. (1999). Multidimensional mechanism design for auctions with externalities. *Journal of economic theory*, 85(2):258–293.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11(1):249–272.

Kang, Z. Y. (2022). Optimal indirect regulation of externalities. Technical report, Stanford University.

Kastl, J., Pagnozzi, M., and Piccolo, S. (2018). Selling information to competitive firms. *RAND Journal of Economics*, 49(1):254–282.

Koessler, F. and Skreta, V. (2022). Informed information design. Technical Report DP17028, CEPR.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge.

Ostrizek, F. and Sartori, E. (2021). Screening while controlling an externality. Technical report, Science Po and University of Naples.

Pai, M. and Strack, P. (2022). Taxing externalities without hurting the poor. Technical report, Rice University and Yale University.

Rodríguez Olivera, R. (2021). Strategic incentives and the optimal sale of information. Technical report, University of Bonn.

Taneva, I. (2019). Information design. *American Economic Journal: Microeconomics*, 11(4):151–85.

Xiang, Y. and Sarvary, M. (2013). Buying and selling information under competition. *Quantitative Marketing and Economics*, 11(3):321–351.