

Selling Information in Competitive Environments*

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

December 15, 2023

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 optimal 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, Roi Orzach, Maryann Rui, and seminar participants at UChicago, UIUC, INFORMS, Harvard, and MIT for insightful comments and valuable feedback.

[†]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* (?). We consider finitely many data buyers and a data seller. The data buyers compete in a simultaneous-move, 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 (?), and the buyers retain control over their downstream actions. Likewise, the

¹We describe the targeted advertising example at length below. See ? and ? for a discussion of direct vs. indirect sales of information.

seller has partial but not full control over each buyer’s outside options. Indeed, the seller cannot prevent any buyer from playing in the downstream game under their prior information only. However, the seller can design the information revealed to any buyer when one or more buyers do not participate in the mechanism, 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 (?).

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 their 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) they are strictly better off following the seller’s recommendation or (b) their opponents do not receive the correct recommendation with probability one. As a result, the seller uses the threat of revealing information to a buyer’s competitors in order to charge a strictly positive price to some types with negative Myersonian virtual values—even types 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 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

²In our approach, we further assume a platform such as Amazon has full commitment power to set information structures. Recent work by ? analyzes the problem of information disclosure to multiple agents by a designer without commitment power.

pay more for targeted campaigns and even more for exclusive access to targeted campaigns.

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, ?? study the sale of information to traders who compete in financial markets. More recently, ? and ? study settings where a single buyer has private information about their 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, ? 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. ? 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., ?, ?, ?, and the references therein. Most papers in this literature view 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, ? study an auction setting with multidimensional and interdependent valuations. Each buyer is privately informed about their valuation for a good and about the externality that they impose 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, ? restrict attention to the second-price auction and

³? 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. ? study the sale of cost information to a large number of perfectly competitive firms, each one facing a privately informed manager. ? consider two sellers who acquire information about a consumer's location along a Hotelling line.

⁴? study the simpler problem where each buyer knows the externality imposed on them by others receiving an object. Under appropriate symmetry assumptions, the problem reduces to a one-dimensional mechanism-design problem. In this vein, ? study a screening model with externalities where each buyer's type affects

study more general externalities in the game played by the buyers. See ? 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 ?. 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 ? and ? 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 among n players (buyers). The game is 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 actions $\mathcal{A}_i$, and a (gross) utility function

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

Information Structures The monopolist information seller chooses a set of signals $\mathcal{S} = \prod_{i=1}^n \mathcal{S}_i$ and a message (bid) space $\mathcal{B} = \prod_{i=1}^n \mathcal{B}_i$, and commits to a communication rule

both their valuation (e.g., for a network good) and the influence their actions (e.g., consumption) impose on other buyers. Their analysis focuses on the countervailing impact of payoff types and influence functions.

$\sigma : \Theta \times \mathcal{B} \rightarrow \Delta(\mathcal{S})$ and payments $p = (p_1, \dots, p_n) \in \prod_{i=1}^n \mathbb{R}_{\geq 0}^{\mathcal{B}_i}$. Given a vector of bids $b \in \mathcal{B}$ and state $\theta \in \Theta$, we write $\sigma(\cdot; \theta, b) : \mathcal{S} \rightarrow [0, 1]$ 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 $u_i(A_i; \theta, V) - p_i(B_i)$.

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 , and therefore we adopt the simpler form $p_i(B_i)$ in the rest of the paper.⁵

The above formulation reduces the problem of information sale to a joint mechanism and information design problem. Throughout this paper and unless specified otherwise, we maintain the following assumptions.

⁵Consider 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 $\tilde{p}_i(v_i) = \mathbb{E}[p_i(B, \theta, R) | B_i]$. In other words, given any feasible mechanism with a general payment of the form $p_i(B, \theta, R)$, the mechanism in which we define $p'_i(B_i) \triangleq \mathbb{E}[p_i(B, \theta, R) | B_i]$ is equivalent from the perspective of player i , satisfies the same constraints, and leads to the same expected revenue.

Assumption 1 (Independent Private Types). *The random variables $(\theta, v_1, \dots, v_n)$ are drawn from a mutually independent prior, and the private types $(v_i)_{i \in [n]}$ are supported on the non-negative reals.*

Assumption 2 (Binary Game with Additive Payoffs). *There are 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*

$$u_i(a; \theta, v) = v_i \cdot \pi_i(a; \theta), \quad (1)$$

where the downstream payoff π_i of buyer $i \in [n]$ is given by

$$\pi_i(a; \theta) = \mathbf{1}\{a_i = \theta\} - \frac{\alpha}{n-1} \sum_{j \neq i} \mathbf{1}\{a_j = \theta\}. \quad (2)$$

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. Furthermore, 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.⁶ 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 $\alpha \geq 0$.⁷

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 $\Theta = \{0, 1\}$ representing an individual consumer's preferences (unknown to the merchants). In this context, the merchants' goal is to match their products 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 (2) for each action profile $a \in \{0, 1\}^2$ in each state of the world θ as

⁶In 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.

⁷The case of positive externalities ($\alpha < 0$) is a straightforward extension of our analysis and is discussed in ??.

	0	1		0	1
0	$1 - \alpha, 1 - \alpha$	$1, -\alpha$	0	$0, 0$	$-\alpha, 1$
1	$-\alpha, 1$	$0, 0$	1	$1, -\alpha$	$1 - \alpha, 1 - \alpha$
	$\theta = 0$			$\theta = 1$	

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, merchant i 's total payoff (gross of any payments to the seller) is given by $v_i \cdot \pi_i$.

First-best Benchmark It will be informative to compare the optimal mechanisms derived in Section 4 to the first-best mechanism that optimizes welfare pointwise for each realization of the state and players' types, ignoring the possibility of misreporting and strategic deviation. For the payoffs (2) we write the welfare as

$$\sum_{i=1}^n v_i \left(\mathbf{1}\{a_i = \theta\} - \frac{\alpha}{n-1} \sum_{j \neq i} \mathbf{1}\{a_j = \theta\} \right) = \sum_{i=1}^n \left(v_i - \frac{\alpha}{n-1} \sum_{j \neq i} v_j \right) \mathbf{1}\{a_i = \theta\},$$

where we changed the order of summation in the second expression. From there, we immediately obtain as the first-best recommendation rule:

$$a_i = \theta \quad \text{if and only if} \quad v_i \geq \frac{\alpha}{n-1} \sum_{j \neq i} v_j. \quad (3)$$

In other words, buyer i receives the “correct” action recommendation ($a_i = \theta$) iff the resulting increase in their utility outweighs the aggregated decrease in other buyers' utilities.

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 ?, 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 space⁸ $\mathcal{V}_i$ 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 the buyer's choice of action by a_i . Incentive compatibility (below) requires each buyer to both report their 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 $\delta : \mathcal{A}_i \rightarrow \mathcal{A}_i$,

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

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

This definition of incentive compatibility is closely related to the one of ?, 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 deviates 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, for each $\delta : \mathcal{A}_i \rightarrow \mathcal{A}_i$ and $v_i \in \mathcal{V}_i$,

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

where A is distributed as $\sigma(\theta, V)$, i.e., 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 $v_i \in \mathcal{V}_i$,

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

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

The first expression shows data buyer i 's strategic behavior before receiving the action recommendation when they intend to report their 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 $\delta : \mathcal{A}_i \rightarrow \mathcal{A}_i$ are more profitable than obedience, i.e., the identity mapping $id : \mathcal{A}_i \rightarrow \mathcal{A}_i$.

⁸When discussing the participation constraint below, we will in fact consider the report space of buyer i to be $\mathcal{V}_i \cup \{\perp\}$ where the additional symbol ' $\perp$ ' is used to indicate the choice not to participate.

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 their 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 ??.

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 $(v_i, v'_i) \in \mathcal{V}_i^2$,

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

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 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 or 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 $\mathcal{B}_i := \mathcal{V}_i \cup \{\perp\}$ for the bid space of buyer $i \in [n]$ and $\mathcal{B} := \prod_{i=1}^n \mathcal{B}_i$, the communication rule is now a function $\sigma : \Theta \times \mathcal{B} \rightarrow \Delta(\mathcal{A})$ with the constraint that it only sends a recommendation to the participating buyers. In other words,

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

Similarly, the payment function $p_i : \mathcal{B}_i \rightarrow \mathbb{R}_{\geq 0}$ of buyer $i \in [n]$ satisfies $p_i(\perp) = 0$.

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

is given by,

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

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

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

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 $i \in [n]$,

$$\mathbb{E}[u_i(A; \theta, V) - p_i(V_i) \mid V_i = v_i] \geq \max_{a_i \in \mathcal{A}_i} \mathbb{E}[u_i(a_i, A_{-i}; \theta, V) \mid V_i = v_i], \quad (4)$$

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 (4). In other words, the seller “punishes” a non-participating buyer by sending optimal recommendations to the remaining buyers 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 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). As discussed above, incentive compatibility rules out double deviations and implies both truthfulness and obedience. Proposition 1 below shows that 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. Note that this converse implication is not true in general but holds here due to our assumption that the buyers' utilities are multiplicatively separable in their independent private types and the outcome of the downstream game (see Eq. 1).

Proposition 1 (Incentive Compatibility Characterization). *A mechanism is incentive compatible whenever it is truthful and obedient.*

Proof. See ??.

□

Truthfulness To characterize truthful mechanisms, we follow the classical result of ?, 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*

$$p_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot \tilde{\pi}_i(\underline{v}) + p_i(\underline{v}) - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds. \quad (5)$$

Proof. See ??.

□

Obedience For the additive payoffs (2),⁹ 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

$$\mathbb{P}[A_i = \theta \mid V_i] = \max_{k \in \Theta} \mathbb{P}[\theta = k] =: P_{\max}. \quad (6)$$

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

Proposition 3 (Obedience Characterization). *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 ??.

□

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 to social welfare and revenue maximization. We show below that, for the additive payoffs in Eq. (2), both objectives can be written as a weighted sum of the probabilities that the mechanism recommends the dominant strategy to each buyer (see Eq. (7) 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 the mechanisms that maximize social welfare and revenue, respectively.

⁹Our characterization does not 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

$$W := \mathbb{E} \left[\sum_{i=1}^n w_i(V) \mathbf{1}\{A_i = \theta\} \right] = \sum_{i=1}^n \mathbb{E} \left[w_i(V) \mathbb{P}[A_i = \theta \mid V] \right] \quad (7)$$

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

Expression 7 and the characterization of obedience obtained in Proposition 3 suggest a convenient parametrization of the seller's problem in terms of the functions $h_i : \mathcal{V} \rightarrow [0, 1]$ given by $h_i(V) := \mathbb{P}[A_i = \theta \mid V]$ 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 that 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 the 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 $\sigma : \Theta \times \mathcal{V} \rightarrow \Delta(\mathcal{A})$ such that almost surely, $\mathbb{P}[A_i = \theta \mid V] = h_i(V)$ for $i \in [n]$.*

Proof. See ?? □

We now describe a general recommendation rule that optimizes criteria of the form (7), which include social welfare and seller revenue, subject to the obedience constraints. Recall the definition of $P_{\max}$ given in (6).

Proposition 5 (Optimal Mechanism). *Consider an objective W of the form (7) where for $i \in [n]$, $w_i : \mathcal{V} \rightarrow \mathbb{R}$ 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 $i \in [n]$, there exists a function $t_i^* : \mathcal{V}_i \rightarrow \mathbb{R}$ such that for all $v_i \in \mathcal{V}_i$,*

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

Then the deterministic recommendation rule given by

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

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

Proof. See ??.

□

To gain intuition for the characterization of optimal mechanisms, note that the objective function W in (7) and the obedience constraints are separable. In other words, the optimization problem reduces to solving separately for each $i \in [n]$ and $v_i \in \mathcal{V}_i$:

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

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 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 the constraint is satisfied. This set corresponds to the level $t_i^*(v_i)$ 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 additive payoffs (2), we can write the expected social welfare as

$$\begin{aligned} W &= \sum_{i=1}^n \mathbb{E} \left[V_i \left(\mathbf{1}\{A_i = \theta\} - \frac{\alpha}{n-1} \sum_{j \neq i} \mathbf{1}\{A_j = \theta\} \right) \right] \\ &= \sum_{i=1}^n \mathbb{E} \left[\left(V_i - \frac{\alpha}{n-1} \sum_{j \neq i} V_j \right) \mathbf{1}\{A_i = \theta\} \right]. \end{aligned} \tag{8}$$

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

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




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 $v^* = F^{-1}(1/2)$ is the median of the type distribution—chosen to be a standard exponential here).

which is of the form (7). We can thus apply Proposition 5 and obtain the following characterization of the welfare-maximizing (second best) mechanism.

Proposition 6 (Welfare Optimal Mechanism). *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. variables¹⁰ distributed according to F . Define $v^* \in \mathbb{R}$ such that*

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

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

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

Proof. See ??.

□

¹⁰ $F^{(k)}$ 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 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 $v_j \leq v^*$. For this rule, the recommendation to buyer i is independent of their type and satisfies $\mathbb{P}[A_i = \theta \mid V_i] = F(v^*) = P_{\max}$. 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 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.¹¹ 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 their type satisfies $v_i \geq \alpha v_j$. This is simply the first-best benchmark (in the absence of the obedience and truthfulness constraints) derived in (3): a buyer is recommended the correct action if their value exceeds the externality they impose on the other buyer. In particular, when buyer i ’s type is large enough compared to buyer j ’s type ($v_i/v_j \geq \max\{\alpha, 1/\alpha\}$), they are recommended the right action exclusively, hence maximizing their 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 efficient allocation recommends the correct action to both buyers. In contrast, when $\alpha > 1$, the region is defined by $v_j/\alpha \leq v_i \leq \alpha v_j$, 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 a collaborative strategy in which both buyers pick the “wrong” action.

¹¹This 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 $\alpha = 3/2$.

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 $v_j \leq v^*$. Distorting buyer i 's recommendation is required, and hence obedience is binding when $v_j \leq v^*$ and $v_i \leq \alpha v_j$.

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, the interim expected payoff of buyer $i \in [n]$, $\tilde{\pi}_i(V_i) := \mathbb{E}[\pi_i(A; \theta) | V_i]$, satisfies almost surely*

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

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

Proof. See ??.

□

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.

As we see from Proposition 7, the truthfulness constraint is not binding in the welfare-optimal mechanism. In other words, the second best mechanism that maximizes welfare

subject only to the obedience constraint satisfies truthfulness “for free.” The distortion that this mechanism introduces compared to the first-best benchmark (3) (namely, that a buyer always receives the good when the sum of their competitors’ types is less than v^*) is solely for the sake of guaranteeing obedience, and no further inefficiency is required to incentivize truth-telling.

To confirm that the inefficiency of the welfare-optimal mechanism is solely due to obedience, it is also easy to verify that the interim payoff $\tilde{\pi}_i$ resulting from the first-best allocation (3) is non-decreasing, implying that without the obedience constraint, it is possible to truthfully implement the first-best mechanism.

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* $\phi : \mathcal{V}_i \rightarrow \mathbb{R}$ defined by

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

is non-decreasing, that is, F is *regular* in the sense of ?.

We first show in Lemma 8 that maximizing the seller’s expected revenue reduces to maximizing the virtual surplus, as in ?.

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 buyer¹² 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 (5) 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*

$$p_i(v_i) = v_i \cdot \tilde{\pi}_i(v_i) - \underline{v} \cdot K - \int_{\underline{v}}^{v_i} \tilde{\pi}_i(s) ds. \quad (9)$$

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

¹²Using 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.

Proof. See ??.

□

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

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

This objective function is of the form (7) 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 (9).

Proposition 9 (Revenue Optimal Mechanism). *Denote by F_ϕ the c.d.f.¹³ 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 $F_\phi^{(n-1)}(\phi(v^\star)) = P_{\max}$ and $\bar{\alpha} := \alpha/(n-1)$. Then, the recommendation rule maximizing virtual surplus subject to obedience is the deterministic rule given by

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

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 $w_i(v) = \phi(v_i) - \bar{\alpha} \sum_{j \neq i} \phi(v_j)$. □

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 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 the same mechanism as in Figure 2 (left) guaranteeing the obedience of buyer i .

¹³When F is a regular distribution, the virtual value function ϕ is invertible.




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 $v^* = F^{-1}(1/2) = \ln 2 < v_0$.

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 ?? 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^* . This in turn depends on the magnitude of the parameter $P_{\max}$ and is discussed in ?? below.

Proposition 10 below gives an expression for the expected downstream payoff $\tilde{\pi}_i$ of each buyer in the obedient mechanism described above. Because ϕ is 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 turn 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 of K , the downstream payoff of a non-participating buyer, we must therefore design the outside option to minimize K .

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, the interim downstream payoff $\tilde{\pi}_i$ of buyer $i \in [n]$ is the non-decreasing function*

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

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 ??.

□

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 $(n-1)\bar{\alpha} = \alpha$, 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 $v^* < v_0$, 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, $F(\phi^{-1}(\phi(v_i)/\alpha)) > F(v^*)$. Buyer i then has a strict incentive to follow the seller's recommendation, i.e., their obedience constraint is slack.
3. Some types of buyer i such that $v^* < v_i$, 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 $[0, v^*]$ 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.

To understand the role of the truthfulness and obedience constraints in limiting the seller's revenue, we now compare the mechanism of Proposition 9 to the benchmark cases in which these constraints are relaxed one at a time.

We first relax the obedience constraint. Observe from the characterization of Proposition 3 that relaxing the obedience constraint is equivalent to formally setting $P_{\max} = 0$. We thus obtain from Proposition 9 that the revenue-optimal recommendation with private information and no obedience is

$$A_i = \theta \quad \text{if and only if} \quad \bar{\alpha} \sum_{j \neq i} \phi(v_j) \leq \phi(v_i). \quad (10)$$

The associated payment is described in Proposition 10, where the critical type v^* is now given by $\phi(v^*) = \underline{v}$. Compared to the recommendation rule in Proposition 9, the seller is no longer required to send the correct action recommendation to buyer $i \in [n]$ when $\sum_{j \neq i} \phi(v_j) \leq$

$\phi(v^*)$. Indeed, this distortion’s purpose was to guarantee that buyer i receives the correct recommendation with probability at least $P_{\max}$ and thus be obedient. Observe also that (10) has the same form as the first-best recommendation rule (3) but with the buyers’ types replaced with their virtual counterparts. In other words, the mechanism takes the form of virtual surplus maximization due to the form of payments imposed by truthfulness.

When we relax truthfulness (i.e., the buyers have no private information) but maintain obedience, the only constraint on payments is the participation constraint. The seller can thus extract the totality of the difference between a buyer’s interim utility and their reservation utility $v_i(P_{\max} - \alpha)$. The problem of maximizing revenue therefore reduces to maximizing welfare, and the optimal allocation in this case is the one given by Proposition 6.

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 information structure and in particular, the buyers’ prior information discussed in ??, the competition structure in the downstream game as captured by the externality parameter α (??), and the number of buyers (??).

5.1 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* 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 two critical aspects of our seller’s problem: the buyer’s reservation




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 $v^* = F^{-1}(1/2) = \ln 2 < v_0$. The third row is the second-best mechanism with an asymmetric prior ($p_{\max} = 3/4$), for which $v^* = \ln 4 > v_0$.