

Selling Information in Games with Externalities

Thomas Falconer¹ Anubhav Ratha² Jalal Kazempour¹ Pierre Pinson^{3,1} Maryam Kamgarpour⁴

Abstract

A competitive market is modeled as a game of incomplete information. One player observes some payoff-relevant state and can sell (possibly noisy) messages thereof to the other, whose willingness to pay is contingent on their own beliefs. We frame the decision of what information to sell, and at what price, as a product versioning problem. The optimal menu screens buyer types to maximize profit, which is the payment minus the externality induced by selling information to a competitor; that is, the cost of refining a competitor's beliefs. For a class of games with binary actions and states, we derive the following insights: (i) payments are necessary to provide incentives for information sharing amongst competing firms; (ii) the optimal menu benefits both the buyer and the seller; (iii) the seller cannot steer the buyer's actions at the expense of social welfare; (iv) as such, as competition grows fiercer it can be optimal to sell no information at all.

1. Introduction

Two key trends mark the rise of today's digital economy: (i) the collection of vast amounts of data from our ever-more digitized lives; and (ii) the advancement of computing algorithms, hardware, and platforms that can transform this data into actionable insights. Many firms are therefore striving to deploy state-of-the-art machine learning models to improve their value propositions. Consequently, the demand for data is growing at an unprecedented rate, to the extent that it has even been dubbed the world's most valuable commodity—the “oil” of the digital age—by The Economist (2017). This raises an important question: *If data is a commodity, then how should we value it?*

Unlike material commodities, data is an *intangible* asset, the very definition of which is a matter of debate. In fact, data is often viewed as an ephemeral entity that can be processed into subjective information, so its value is not an intrinsic property, but dependent on when, how, and by whom it is

actually used (Floridi, 2002; 2009). We adopt this view by modeling a *buyer* as a decision-maker under uncertainty that seeks additional information to refine their beliefs about some state of the world. In turn, a monopolistic *seller* owns data relevant to the buyer's decision, and offers (possibly noisy) messages thereof in exchange for money, for who the value of this information depends on their prior beliefs.

We explore what information, and at what price, the seller should offer to maximize profit. We assume that both parties compete in a downstream market, so the seller has information relevant to each of their payoffs. This is motivated by growing concerns over the relationship between data and the distribution of market power in many industries, insofar that competing firms face asymmetric access to information. For example, in June 2021, the EU launched an antitrust investigation into whether Google distorted competition by restricting third party access to user information for advertising, reserving exclusive use for itself (European Commission, 2021). Central to this debate is whether firms should share information to benefit social welfare, or whether they can sell their information for profit.

However, packaging and pricing information to monetize it within competitive environments is not a simple task, as the seller needs to consider not only the buyer's utility, but also any externalities, both positive or negative, they may induce by sharing information with competitor. We model a competitive market environment as a game of incomplete information, where the players face a common *fundamental uncertainty* associated with a payoff-relevant state of the world. Each player also observes a *private* signal that determines their beliefs about the fundamental uncertainty, meaning they are also subject to *strategic uncertainty* with respect to the signals received by the other players. Players then simultaneously select an action to maximize their expected payoffs. We extend this framework by including an interim step where one player (the seller) observes the realized state before actions are chosen.

The buyer's beliefs, and thus their valuation for information, is unknown to the seller, who designs and offers a menu of communication rules, which are likelihood functions which prescribe a distribution over messages the buyer could receive from the seller conditioned on the realized state. It can therefore be characterized by a degree of obfuscation of the

¹DTU, DK ²Vestas, DK ³Imperial College London, UK ⁴EPFL, CH. Correspondence to: Thomas Falconer <falco@dtu.dk>.

true data. This setup is adapted from those seen in information design literature, where a social planner commits to a communication rule to influence the behavior of players in a game (Bergemann & Morris, 2019). This is equivalent to selecting the Bayes correlated equilibrium that maximizes the planner’s objective (Bergemann & Morris, 2016). In our case, communication rules have an associated price and the problem becomes one of joint information and mechanism design, in which the seller must elicit the buyer’s valuation for information.

1.1. Contributions

Before discussing related works, we summarize our contributions and outline of the paper:

- In Section 2, we formalize a two-player game of incomplete information that we use to model a competitive market. We adopt a simplified setup where both action sets and the state of the world are binary, and consider the case wherein each player has a dominant strategy to match their action with the state. As such, the buyer’s private type is one-dimensional and their valuation is both piecewise linear and independent of the seller’s action. Since the communication rule purchased influences the buyer’s decisions, the seller has (anti-) coordination incentives if their expected payoff (decreases) increases with the probability that the buyer also chooses the correct action. Although we focus on competitive environments with anti-coordination incentives, we show that our setup extends to the opposite case which could be viewed as a game with strategic complements instead of strategic substitutes.
- In Section 3, we outline the aspects that constrain the menu offered by the seller. Based on the principle that information is only valuable insofar as it changes the buyer’s action, the seller’s messages are viewed as action recommendations. We establish in Definitions 3.4 and 3.5 that a menu is incentive compatible if the buyer is best off truthfully reporting their type *and* obediently following the recommended action.
- In Section 4, we adapt the Myersonian auction format to provide novel insights into selling information to a competitor (see Proposition 4.3). To study the nuances of this product versioning (or second-degree price discrimination) problem, we first consider just two buyer types, ranked according to their valuation of the fully informative communication rule (see Corollary 4.4). We show that the seller can reduce the information rent of the less informed type by offering partial information depending on the type congruence—whether or not they would take different actions given their prior beliefs—, which is consistent with previous findings

(Bergemann et al., 2018). However, unlike this prior work, we show that competition between the buyer and the seller induces an externality in the form of a cost of refining a competitor’s beliefs. This makes the optimal menu a function of both the intensity of competition and the seller’s own beliefs. For instance, if types are congruent, such that they would both select the same action without additional information, and the seller’s own beliefs are strongly in the opposite direction, the externality exceeds the maximum revenue and no information is sold. If the seller’s beliefs are in the same direction, they may choose to sacrifice some revenue by offering more information to the lower type to reduce the expected externality.

We show for continuous type spaces that the mechanism design problem reduces to maximizing virtual surplus (see Corollary 4.7). Our results imply monetary compensation can serve as an effective incentive for information sharing, with the seller able to package and price information for profit, even in competitive environments where it would not otherwise occur organically. The profitability is, however, sensitive to the seller’s own beliefs and the intensity of competition, as well the distribution of buyer types, which generally occur with non-zero probability both within and across classes of congruent types. Lastly, as any buyer could ignore the seller’s message, the seller’s cannot design a menu that maximizes profit at the expense of social welfare. This explains why we observe that as competition grows fiercer it can be optimal to sell no information at all. Intuitively, the seller wants to message the buyer to take the action opposite to their own beliefs with certainty. However, such messages would not satisfy Bayesian consistency, meaning that they are not rational under the prior, and the buyer would not follow the recommended action. This result suggests that fiercely competitive environments, where the externality cost renders information sales unprofitable, may require additional incentives.¹

- Finally, in Section 5 we gather a set of conclusions and perspectives for future work.

1.2. Related Works

We now discuss our contributions to the literature on information economics covering three key areas: (i) providing incentives for collaborative analytics; (ii) facilitating information sharing in competitive environments; and (iii) selling information to imperfectly informed decision-makers.

¹The code used to run all of our simulations is publicly available at the following repository: www.github.com/tdfalco/trading-information

Collaborative Analytics. For machine learning tasks, raw datasets are typically acquired directly from their owners via bilateral transactions (Rasouli & Jordan, 2021). However, valuing raw data is challenging because it depends on the information the buyer derives, which is difficult to assess before they actually use it. In addition, the value of information is inherently combinatorial, as datasets invariably contain correlated signals, which also makes it challenging to ensure privacy if one’s information can be inferred from that of others (Acemoglu et al., 2022; Fallah et al., 2024a;c). Thus, these seemingly straightforward transactions can quickly become intractable to price.

To address this challenge, recent works propose *analytics markets*: real-time mechanisms that match datasets to machine learning tasks based on predictive value, without transferring raw data. These markets leverage the privacy benefits of collaborative analytics like federated learning (Zhang et al., 2021), but add monetary incentives to take part (Fallah et al., 2024b). First proposed by Agarwal et al. (2019), analytics markets aggregate features from multiple sellers. Buyers submit tasks and bids reflecting their valuation for improved predictions. The platform determines how much information is sold and at what price, allocating market revenue to sellers according to their contributions to accuracy improvements. Platforms for classification (Koutsopoulos et al., 2015) and regression (Pinson et al., 2022) tasks have been proposed, tackling challenges related to strategic behavior (Falconer et al., 2025), timing of data availability (Feng et al., 2021) and financial security of the sellers (Falconer et al., 2024).

These works challenge the notion that raw data has intrinsic value, arguing that its usefulness depends on the information derived. While this view of data valuation is intuitive, it is typically assumed that buyers truthfully report how much they value predictive accuracy, an assumption difficult to sustain in real competitive environments, where sellers may also have incentives to withhold (Kakhbod et al., 2021) or distort (Ziv, 1993) information to benefit themselves, even at the expense of social welfare. By treating the platform as a monopolistic seller, clearing these markets can be viewed as a product versioning problem, akin to our setup, where the seller strategically designs offerings to maximize profit. Our work thus offers an initial step toward tackling the broader challenge of pricing information in analytics markets.

Competitive Environments. Early research into incomplete information games found that when multiple equilibria exist without strategic uncertainty, introducing information asymmetry can eliminate this multiplicity. In other words, strategic uncertainty can coordinate players towards a particular equilibrium (Hellwig, 2002). This notion underpins Bergemann & Morris (2016)’s concept of Bayes correlated equilibrium, an extension of Aumann (1987)’s complete in-

formation correlated equilibrium to Bayesian games, where a mediator privately recommends actions to players based on the realized uncertainties, in a way that players willingly follow the recommendations. Similar insights have shown that players can strategically share information with select others to achieve self-serving outcomes (Dahleh et al., 2016). This connects to the broader literature on information asymmetry and market power, surveyed in detail by Bergemann & Bonatti (2019). In seminal works, Admati & Pfleiderer (1986; 1990) model a monopolist that sells information to traders, showing that offering noisy signals tailored to each maximizes profits by limiting information leakage through stock prices. Admati & Pfleiderer (1988) build on this by allowing the seller to also trade in the market, finding that the value of selling information increases with buyers’ risk aversion.

Information sharing in incomplete information games, illustrated in Figure 1a, has also been explored in game theory literature; key results are summarized by Raith (1996). Works such as Novshek & Sonnenschein (1982); Clarke (1983); Vives (1984); Gal-Or (1985); Shapiro (1986) model strategic uncertainty in oligopolies, considering the nature of both competition (e.g., strategic complements or substitutes) and fundamental uncertainties (e.g., costs or capacities). A common result is that whilst information sharing can improve production efficiency and social welfare, firms often resist it without compensation, as it may redistribute market power and lower individual profits. In our work, we aim to explore a further dimension: how the intensity of competition, not just its nature, shapes the value of information sharing.

Selling Data. Our work is most closely related to recent applications of mechanism design to provide incentives for information sharing. To maintain tractability, this often requires abstracting both the competitive environment as a Bayesian game. For example, Babaioff et al. (2012) study a setting where buyer utility depends on two uncertain states, with the buyer and seller each privately informed about one. They characterize the optimal mechanism as conditional on the seller’s observed state, allowing the seller to exploit correlations between their private information and the buyer’s type, similar to Crémer & McLean (1988). Giving the seller flexibility to price and allocate information after observing the state parallels classical communication problems such as cheap talk (Crawford & Sobel, 1982), verifiable message (Milgrom, 1981), and signaling games (Spence, 1978). In these settings, a mediator, the seller in our case, chooses what to communicate after observing the state. This leads to an intractably large decision space as receivers may discount the message unless their incentives are sufficiently aligned with the mediator’s.

A simpler approach lends itself to the information design literature (Bergemann & Morris, 2019), where the mediator




Figure 1. Existing frameworks for sharing information. The blue, red and green arrows indicate monetary transactions, strategic actions, and information flows between the firms and the market, respectively. Whilst in each framework, Firm B shares information with Firm A, the setups differ as follows: in (a) information is shared freely between competitors, in (b) information is purchased yet Firm B is a third party that doesn't compete with Firm A, and in (c) information is purchased and both firms compete in the market.

commits *ex-ante* to a communication rule. This commitment dramatically simplifies the analysis and often yields a unique solution. Although such commitment may be unrealistic in practice, the framework is valuable as a benchmark for understanding the range of achievable outcomes under different information structures. This perspective mirrors that of mechanism design, which also assumes commitment, such as in auctions, and remains useful as a conceptual tool even without a literal mediator, as in Myerson & Satterthwaite (1983)'s analysis of bargaining. If the mediator holds no informational advantage over the players, information design collapses to standard communication games (Myerson, 1986). With a single receiver, information design reduces to Bayesian persuasion (Kamenica & Gentzkow, 2011).

Information design literature typically does not study the inclusion of monetary transfers. In Bergemann et al. (2018; 2022) this problem is framed as designing an auction, differing from Babaioff et al. (2012) in requiring the seller to commit to a mechanism *ex-ante*. We adopt this framework, which follows the Bayesian model of decision-making under uncertainty from Blackwell (1951; 1953), where the buyer purchases a communication rule, or a so-called Blackwell experiment, rather than paying for a message revealed *ex-post* by the seller. Unlike standard information design, which assumes a common prior (Taneva, 2019; Mathevet et al., 2020) or an omniscient mediator, we study elicitation of private information and its influence on what information is sold. Cai & Velez (2020) and Ravindranath et al. (2024) build on Bergemann et al. (2018), focusing on efficient algorithms for approximating optimal menus.

These prior works do not consider externalities. One common externality arises when a third party sells information to multiple competing buyers, as illustrated in Figure 1b, making the value of information to one buyer dependent on what others receive. Rodríguez Olivera (2024) extend Bergemann et al. (2018) to a multi-buyer setting, where the optimal menu depends on buyers' strategic behavior and the correlation of their private information. Externalities are also modeled explicitly in Bimpikis et al. (2019), Agarwal et al. (2024), and Bonatti et al. (2024). The latter studies a two-player game with binary actions and state space, akin

to a discrete Hotelling competition (Hotelling, 1929), where firms choose locations without knowing consumer distributions. Here, the seller, informed about the distribution, exploits asymmetries between firms to design optimal communication rules and shows that limiting the set of buyers that receive information can maximize profit.

We focus on a related but orthogonal problem of selling information not *to*, but *amongst* competing firms, illustrated in Figure 1c. Gradwohl & Tennenholtz (2023) study a similar problem, framed as a variant of Hotelling's model where firms compete on price, but only one firm knows consumer locations. As in Raith (1996), they find that whilst full disclosure harms both firms, selling a subset of information can benefit both. Similarly, Castiglioni et al. (2023) model the sale of information to a budget-constrained competitor using information design, characterizing an optimal menu of communication rules via a polynomial-time linear relaxation. However, the buyer holds no private information about the state, which may not be true in practice.

In our more general setup, the buyer refines their *own* private beliefs, rather than simply updating a common prior, using information from the seller. This can also be viewed as a variant of Hotelling's model, where firms compete on location and one firm knows the true customer distribution while the other faces uncertainty. Overall, our contributions to the literature are threefold: (i) we offer a first step toward modeling practical analytics markets; (ii) we characterize incentives for information sharing in competitive environments as a function of competition intensity; and (iii) we generalize Bergemann et al. (2018) to incorporate competition between buyer and seller, adapt the frameworks of Bimpikis et al. (2019), Agarwal et al. (2024), and Bonatti et al. (2024) by embedding the seller within the downstream game, and extend the settings of Gradwohl & Tennenholtz (2023) and Castiglioni et al. (2023) by allowing buyers to hold private information about the state.

2. Model

We model two firms competing in a market as an incomplete information game parameterized by a random variable X

with range $\mathcal{X} \subseteq \mathbb{R}$, where each $x \in \mathcal{X}$ is a realizable state of the world. Each player selects an action $a_i \in \mathcal{A}_i$ and receives *ex-post* utility $u_i : \mathcal{A} \times \mathcal{X} \mapsto \mathbb{R}$, for every $i \in \{b, s\}$, where subscripts b and s are the indices of the buyer and seller, respectively, so $\mathcal{A} = \mathcal{A}_b \times \mathcal{A}_s$. We restrict our analysis to binary games, with states $\mathcal{X} = \{0, 1\}$ and action sets $\mathcal{A}_i = \{0, 1\}$. Each player's utility is assumed to consist of a nonnegative term u_i^+ which depends only on their own action, and a nonpositive term u_i^- which depends only on the action of the other player, such that $u_i(\mathbf{a}, x) = u_i^+(a_i, x) + u_i^-(a_j, x)$. Thus, for every (\mathbf{a}, x) pair, we can subtract a state-dependent constant so that, without loss of generality, player i seeks to match their action a_i to the state x and u_i reduces to

$$u_i(\mathbf{a}, x; \tau) = \mathbf{1}_{\{a_i=x\}} - \tau \mathbf{1}_{\{a_j=x\}}, \quad (1)$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function, returning 1 if the statement $\{\cdot\}$ is true, and 0 otherwise. The parameter $\tau \in \mathbb{R}_{\geq 0}$ encodes the intensity of competition between the two players, reducing the payoff if the other player also chooses the correct action.

As X is binary, the common prior is a Bernoulli distribution parameterized by $v \in \mathbb{R}_{[0,1]}$, with

$$p(x; v) = x(1 - v) + (1 - x)v,$$

with v the prior probability $P(X = 0)$.²

Private Information. Each player also receives a private signal $s_i \in \mathcal{S}_i$, with $s_i \sim p(s_i|x)$ its distribution given the state. We assume that private information stems from independent sources, such that $p(s_b, s_s|x) = p(s_b|x)p(s_s|x)$, a common assumption in analysis of incomplete information games. Each player uses their private signals together with the common prior to form updated beliefs via Bayes' rule. With a Bernoulli prior, the posterior retains the same form for any likelihood, so player i 's posterior is

$$p(x|s_i; v_i) = \frac{p(s_i|x)p(x; v)}{\sum_{x' \in \mathcal{X}} p(s_i|x')p(x'; v)} \quad (2)$$

$$= x(1 - v_i) + (1 - x)v_i,$$

with $v_i \in \mathcal{V}_i$ player i 's posterior probability $P(X = 0)$, where $\mathcal{V}_i = \mathbb{R}_{[0,1]}$. Assuming the likelihood is well-defined and non-zero for at least one $x \in \mathcal{X}$, a proper prior ensures a proper posterior and that the posterior update is a martingale. In other words, the expectation of the posterior belief, taken over the marginal distribution of player i 's private signal s_i , equals the prior belief, that

²For a random variable X , its cumulative distribution function is $F(x) = P(X \leq x)$. If X is continuous, it has a density function $p(x)$, satisfying $F(x) = \int_{-\infty}^x p(z)dz$, and if X is discrete, p is a mass function and $F(x) = \sum_{z \leq x} p(z)$.

is, $\mathbb{E}_{s_i \sim p(s_i)}[p(x|s_i; v_i)] = p(x; v)$.³ Hence, the posterior assigns positive probability only to states that had positive prior probability, i.e., no state is plausible if it was ruled out *a priori*. For further intuition about the role of private information, with competition parameterized by τ , we provide the following example.

Illustrative Example. Consider a variant of Hotelling's spatial competition model where, instead of customers being uniformly distributed along an interval, all are concentrated at one of two locations, $x \in \{0, 1\}$. Two competing firms must choose a location $a_i \in \{0, 1\}$ (interpreted as selecting what or where to produce). Firm i knows the true location and always chooses $a_i = x$. If $X = 1$ and firm j chooses incorrectly, which occurs when $v_j \geq 1/2$, then firm i captures the entire market, so $u_i = 1$ and $u_j = 0$. If instead $v_j < 1/2$, firm j also chooses the correct location, leading to competition: both firms split the market, with each earning $u_i = u_j = 1 - \tau$. The parameter τ captures the intensity of competition; as τ increases, competition reduces profits, culminating in a zero-sum game at $\tau = 1$. For $\tau > 1$, the cost of competition outweighs the benefits of choosing the correct action, creating a prisoners' dilemma.

Remark 2.1. Even though we treat τ as constant, our setup readily extends to cases where $\tau : \mathcal{X} \mapsto \mathbb{R}_{\geq 0}$ is a function of the state (e.g., differentiated products) or a function of the other player's action. One may also consider players to have coordination incentives, with $\tau < 0$. In this case, each player would benefit from both themselves and the other player choosing the action that matches the state.

The Seller's Problem. We further assume that before the players choose an action, the seller observes an additional signal that reveals the true state. The seller can then offer (possibly noisy) messages to the buyer in exchange for money. In game-theoretic terms, the seller acts as a mediator, sending messages $m \in \mathcal{M}$ conditioned on the state.⁴ However, unlike in standard information design, the seller is not omniscient, as they do not know the buyer's private information. This setup naturally leads to an information design problem with elicitation, as the seller must first elicit a report from the buyer and then tailor their message. As we include monetary transfers, we frame this problem as one of joint information and mechanism design.

From a mechanism design view, a player's posterior belief $v_i \in \Delta(\mathcal{X})$ serves as their *type*, which can be viewed as a realization of a random variable V_i . The common prior and

³A discrete-time *martingale* is a stochastic process (i.e., a sequence of random variables indexed by time t), $\{X_t\}_{t \geq 1}$, that satisfies the following conditions: $\mathbb{E}[|X_t|] < \infty$ for every $t \geq 1$; and $\mathbb{E}[X_{t+1}|X_1, \dots, X_t] = X_t$ for every $t \geq 1$.

⁴It's not essential that the seller observes the true state; it suffices that they can send messages correlated with it.

the distribution over signals together induce a distribution over interim beliefs $p(v_i) \in \Delta(\Delta(\mathcal{X}))$, a distribution over distributions, determined by the topology of the set of priors, $\Delta(\mathcal{X})$, and the set of distributions over signals, $\Delta(\mathcal{S}_i)$. The buyer seeks to refine their beliefs by acquiring additional information from the seller, with their valuation depending on their private type v_b which is unknown to the seller. To elicit this information, the seller asks the buyer to submit a bid $b \in \mathcal{B}$. Let M denote a discrete random variable with support \mathcal{M} . The seller commits to an *ex-ante* communication rule $p(m|x; b) \in \Delta(\mathcal{M})$, specifying a distribution over messages conditional on the state x for each possible bid b . Each communication rule is paired with a price, determined by a transfer function $t : \mathcal{B} \mapsto \mathbb{R}_{\geq 0}$.

Upon receiving message $m \in \mathcal{M}$, the buyer again updates their beliefs via Bayes' rule, with posterior

$$\begin{aligned} p(x|m, s_b) &= \frac{p(m|x; b)p(x|s_b; v_b)}{\sum_{x' \in \mathcal{X}} p(m|x'; b)p(x'|s_b; v_b)} \\ &= x(1 - \theta_m^b) + (1 - x)\theta_m^b, \end{aligned}$$

which again is a Bernoulli distribution, but now parameterized by $\theta_m^b \in \mathcal{V}_b$, the updated posterior probability $P(X = 0)$ of the buyer upon receiving message m provided they bid b .

Timing. Viewing the interaction between the buyer and the seller as an extensive form game, the timing is:

1. Each player observes their private type $v_i \in \mathbb{R}_{[0,1]}$.
2. The seller offers a menu of communication rules and the buyer reports bid $b \in \mathcal{B}$.
3. State $x \in \{0, 1\}$ is realized and then the buyer receives message $m \sim p(m|x; b)$.
4. The buyer updates their beliefs to θ_m^b and pays $t(b)$ to the seller.
5. Both players choose an action $a_i \in \{0, 1\}$ and obtain *ex-post* utility $u_i(a, x)$.

Equilibrium. The optimal strategy $\sigma : \mathcal{V}_i \mapsto \mathcal{A}_i$ defines a mapping from types to actions. Given the utility function (1) this is simply to pick the action that matches the state with the largest probability mass, such that

$$\sigma(z) = \arg \max_{x \in \{0,1\}} p(x|s_i; z) = \mathbf{1}_{\{z < \frac{1}{2}\}}, \quad (3)$$

for every $z \in \mathcal{V}_i$. We use z to denote the type here since the buyer's may be their prior v_b or θ_m^b after receiving message m . As σ depends only on each player's own type, neither

can improve their *ex-post* utility by unilaterally deviating. Thus, (3) is a dominant strategy, such that

$$u_i(\sigma(z), \sigma(v_s), x; \tau) \geq u_i(a, x; \tau),$$

for every $a \in \mathcal{A}$, $z \in \mathcal{V}_b$, and $v_s \in \mathcal{V}_s$. Of course, as the seller observes the state, their optimal strategy is $\sigma : \mathcal{X} \mapsto \mathcal{A}_s$, where $\sigma(x) = x$. However, this is simply a form of (3) after updating beliefs to $P(x) = 1$.

Remark 2.2. From the perspective of information design, we have a single mediator and a single receiver, and are thus in a one-player Bayesian Persuasion case as in Kamenica (2019), and a communication rule reduces to a communication rule in the sense of Blackwell (1951; 1953).

In summary, the buyer seeks to refine their beliefs by acquiring additional information from the seller. The seller's goal is to design a mechanism, a menu of communication rules and corresponding prices, that maximizes expected profit. In the following section, we outline the constraints that limit the space of feasible menus available to the seller.

3. Mechanism Characteristics

To determine the space of feasible menus, we first characterize the buyer's valuation of a given communication rule, framed as the *gain* in expected utility with and without the additional information from the seller. This characterization allows us to determine under which conditions the buyer has an incentive to participate and to truthfully report their type. The latter incentive is crucial: reasoning over a generic bid space \mathcal{B} is intractable due to the wide range of possible strategic behaviors, but the revelation principle of Myerson (1981) allows us to consider only direct mechanisms where buyers report types truthfully. Thus, we may assume $\mathcal{B} = \mathcal{V}_b$, provided the mechanism is incentive compatible. In other words, bid b can be viewed as a realization of a V_b -measurable random variable, allowing the mechanism to infer the buyer's true type v_b .

Value of Communication Rules In order to quantify the value of a communication rule for type v_b , we first derive the buyer's expected utility using only their private information as follows:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}, \mathcal{V}_s} [u_b(a, x; \tau)] &= \sum_{x \in \mathcal{X}} p(x|s_b; v_b) \int_{\mathcal{V}_s} p(v_s) u_b(a, x; \tau) dv_s \\ &= \sum_{x \in \mathcal{X}} p(x|s_b; v_b) \left(u_b^+(\sigma(v_b), x) \right. \\ &\quad \left. + \int_{\mathcal{V}_s} p(v_s) u_b^-(\sigma(v_s), x; \tau) dv_s \right) \\ &= \mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(v_b), x)] + \mathbb{E}_{\mathcal{X}, \mathcal{V}_s} [u_b^-(\sigma(v_s), x; \tau)], \end{aligned}$$

where given the separability of (1), only the nonpositive part of the utility depends on uncertainty in the seller's private information v_s . After purchasing a communication rule by reporting bid b , the buyer's expected utility conditioned on message m is given by

$$\mathbb{E}_{\mathcal{M}, \mathcal{V}_s} [u_b(\mathbf{a}, x; \tau) | m] = \mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(\theta_m^b), x) | m] + \mathbb{E}_{\mathcal{X}, \mathcal{V}_s} [u_b^-(\sigma(v_s), x; \tau)],$$

hence, the expected utility given a particular communication rule is calculated by integrating over the set of all possible messages, such that

$$\begin{aligned} & \mathbb{E}_{\mathcal{M}, \mathcal{X}, \mathcal{V}_s} [u_b(\mathbf{a}, x; \tau)] \\ &= \sum_{m \in \mathcal{M}} p(m; b, v_b) \mathbb{E}_{\mathcal{X}, \mathcal{V}_s} [u_b(\mathbf{a}, x; \tau) | m] \\ &= \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] + \mathbb{E}_{\mathcal{X}, \mathcal{V}_s} [u_b^-(\sigma(v_s), x; \tau)], \end{aligned}$$

which, in theory, could be intractable if the message space \mathcal{M} is unbounded. However, we now define a special subset of communication rules that limits its cardinality.

Definition 3.1 (Direct Communication Rule). A communication rule is *direct* if every message with positive probability leads to a different optimal choice of action.

Proposition 3.2. *Every communication rule induced by the optimal menu can be direct.*

Proof. Provided in Appendix A. \square

A result of Proposition 3.2 is that it suffices for the seller to only consider direct communication rules where \mathcal{M} has cardinality equal to that of \mathcal{A}_b , in our case, $|\mathcal{M}| = 2$. The proof follows the insight of Blackwell (1951; 1953) that information is only valuable insofar that it changes the receiver's action. If $\mathcal{M} = \{m_0, m_1\}$, a communication rule is two probabilities: $P(m_0 | X = 0; b)$ and $P(m_1 | X = 1; b)$, as illustrated in Table 1.

Table 1. Tabular illustration of a direct binary communication rule.

	m_0	m_1
$X = 0$	$P(m_0 X = 0; b)$	$1 - P(m_0 X = 0; b)$
$X = 1$	$1 - P(m_1 X = 1; b)$	$P(m_1 X = 1; b)$

Given the ultimate aim to match one's action with the state, it is natural to order messages so that m_0 and m_1 recommend actions $a_b = 0$ and $a_b = 1$, respectively. This is achieved by ensuring $\theta_{m_0}^b \geq 1 - \theta_{m_1}^b$, as recall the dominant strategy is to choose the action corresponding to the state with largest probability mass. In this case,

$$P(m_0 | X = 0; b)b \geq P(m_0 | X = 1; b)(1 - b), \quad (4)$$




Figure 2. One-dimensional informativeness.

and

$$P(m_1 | X = 0; b)b \leq P(m_1 | X = 1; b)(1 - b), \quad (5)$$

thereby inducing the respective actions.⁵ Lastly, given the parametrization in Table 1, any communication rule can be fully characterized by its *informativeness*, denoted $I(b) \in \mathbb{R}_{[-1, 1]}$, defined as

$$I(b) = P(m_0 | X = 0; b) - P(m_1 | X = 1; b),$$

which we illustrate in Figure 2. The center, $I(b) = 0$, is the fully informative communication rule that invariably reveals the true state. In each half, one state is revealed with certainty, and the other gets increasingly obfuscated towards the boundary. The boundary identifies extreme points where one message occurs with probability 1, which we later show is the same as revealing no information to the buyer.

We define the gain δ of communication rule I for a certain type v_b provided they bid b as the difference between the buyer's expected utility with and without additional information from the seller such that

$$\begin{aligned} \delta(I(b), v_b) &= \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] - \mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(v_b), x)], \end{aligned} \quad (6)$$

which is determined by the quality of the buyer's private information. Lastly, we assume that v_b provides no additional information about the distribution of messages conditioned on the state, i.e., the buyer's private signal $s_b \in \mathcal{S}_b$ and the message received from the seller $m \in \mathcal{M}$ are independent conditioned on the state, meaning they draw their information from independent sources.

In summary, we have reduced the mechanism design problem to finding a menu consisting of: (i) a single-dimensional allocation function, namely, the informativeness of the messages given the bid, $I : \mathcal{B} \rightarrow \mathbb{R}_{[-1, 1]}$; and (ii) a transfer function $t : \mathcal{B} \rightarrow \mathbb{R}$, specifying how much the buyer should pay for this information.

Feasible Menus. The use of monetary transfers leads to several key distinctions compared to pure information design. First, the buyer's willingness to participate depends not only on the gain of the purchased communication rule but also on its price.

⁵Inequalities (4) and (5) compare unnormalized posteriors, but since both sides of the inequality share the same denominator, we need not write it explicitly.

Definition 3.3 (Individual Rationality). A mechanism is *individually rational* if the buyer is guaranteed at least as much expected utility as not participating, provided they report their true type, which, for every $v_b \in \mathcal{V}_b$, implies

$$\delta(I(v_b), v_b) \geq t(v_b).$$

The transfers also affect incentive compatibility.

Definition 3.4 (Truthfulness). A mechanism is *truthful* if, assuming it is *obedient* as defined below, for every $v_b \in \mathcal{V}_b$ and $b \in \mathcal{B}$, it holds that

$$\delta(I(v_b), v_b) - t(v_b) \geq \delta(I(b), v_b) - t(b).$$

Each truthfully reported type also has to be willing to follow the recommended action.

Definition 3.5 (Obedience). A mechanism is *obedient* if the buyer's best response is the seller's action recommendation, provided they report their type truthfully, such that for every $a_b \in \mathcal{A}_b$ and $v_b \in \mathcal{V}_b$, it holds that

$$\mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^{v_b}), x)] \geq \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(a_b, x)].$$

Any obedient mechanism induces a distribution of actions that is Bayes correlated equilibrium (Bergemann & Morris, 2016). Both truthfulness and obedience together make the mechanism robust to double-deviations, where the buyer misreports their private type *and* deviates from the seller's recommendation. Therefore, together these constraints form incentive compatibility.

Definition 3.6 (Incentive Compatibility). A mechanism is *incentive compatible* if

$$\begin{aligned} \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^{v_b}), x)] - t(v_b) \\ \geq \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(a_b, x)] - t(b), \end{aligned}$$

for every $a_b \in \mathcal{A}_b$, $v_b \in \mathcal{V}_b$ and $b \in \mathcal{B}$.

Satisfying these individual rationality and incentive compatibility constraints limits the space of feasible menus available to the seller. In the next section, we formulate the seller's objective and incorporate these constraints to fully characterize the optimal mechanism design problem.

4. Profit-Maximizing Mechanism

To derive the seller's objective, we first return to the value of a communication rule to simplify the expressions for the buyer's gain and the seller's cost. This allows us to finally design and analyze the optimal mechanism, first considering a simpler setting with two buyer types before generalizing to a continuous type space.




Figure 3. Feasible region of $I_0(b)$ and $I_1(b)$. The ceiling of the feasible region is highlighted in red.

The Buyer's Gain. With the dominant strategy in (3), the expected value of the nonnegative term of the buyer's utility function before any information sharing reduces to

$$\mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(v_b), x)] = (v_b \vee (1 - v_b)), \quad (7)$$

and similarly, after purchasing a communication rule,

$$\begin{aligned} \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] = \\ \sum_{m \in \mathcal{M}} p(m; b, v_b) (\theta_m^b \vee (1 - \theta_m^b)). \end{aligned}$$

where $a \vee b = \max(a, b)$. Next, we use these expressions to simplify the characterization of the informativeness of a communication rule.

Proposition 4.1. *Every communication rule in the optimal menu will reveal at least one of the states with certainty.*

Proof. Provided in Appendix B. \square

To provide intuition for this proof let $I_0(b) = P(m_0|X = 0; b)$ and $I_1(b) = P(m_1|X = 1; b)$ fully define a communication rule as in Table 1, with $I_0(b), I_1(b) \in \mathbb{R}_{[0,1]}$. By Proposition 3.2, $I_0(b) + I_1(b) \geq 1$, for which the feasible region is plotted in Figure 3. It follows that in the optimal mechanism, all feasible objective values with corresponding solution (I_0^*, I_1^*) define a level set which contains a solution (I'_0, I'_1) on the ceiling of the feasible region, on which either measure is always 1, revealing the corresponding state with certainty. It is this line that describes the informativeness $I(b) \in \mathbb{R}_{[-1,1]}$ of the communication rule.

Corollary 4.2. *From Proposition 4.1, the gain δ of a communication rule $I(b)$ to type v_b provided they report b as described in (6) can be written as*

$$\begin{aligned} \delta(I(b), v_b) = 1 - (v_b \vee (1 - v_b)) \\ - I(b) (\mathbf{1}_{\{I(b) \geq 0\}} - v_b). \end{aligned} \quad (8)$$




Figure 4. Gain as a function and buyer types, as described in (8), for $I = 0$ (solid), $I = 1/2$ (dashed) and $I = -1/2$ (dotted).

This expression matches that for the buyer's valuation in Bergemann et al. (2018), as dominant strategies ensure that selling information to a competitor only affects u_b^+ and u_s^- , the latter of which we later show only appears in the objective function. To remain self-contained, Figure 4 shows the gains from several communication rules, illustrating how the value of information varies with the buyer's type. Observe that: (i) the fully informative communication rule is the most valuable for every type, with $\delta(I, v_b) = 1/2$; (ii) the type with the highest willingness to pay is that which is least informed (i.e., $v_b = 1/2$) and those with the lowest willingness to pay are those that are most informed (i.e., $v_b \in \{0, 1\}$); and (iii) the distance $|v_b - 1/2|$ alone is not a sufficient statistic for the value of information, as asymmetries arise based on interim beliefs.

These observations imply that these information commodities are inherently two-dimensional, comprising both *quality* and *value*; accuracy is only valuable so long as it changes the buyer's action. This distinction has been well studied in the field of predictive analytics, including in meteorological forecasting (Murphy, 1993). It implies that a direct communication rule yields a nonnegative gain: the message either (i) changes the buyer's best response, leading to a weak utility increase; or (ii) leaves the action unchanged, resulting in zero value. *Hereafter, we consider only truthful mechanisms with $b = v_b$.* We hence obtain: $-1 \leq I(v_b) \leq v_b/(1 - v_b)$ if $v_b \leq 1/2$ and $-(1 - v_b)/v_b \leq I(v_b) \leq 1$ if $v_b > 1/2$. Hence, communication rules $I(v_b) \in \{-1, 1\}$ indeed reveal no information, as the types able to receive these would already take the recommended action without additional information.

The Seller's Cost. For a given communication rule, the seller's cost is the expected externality incurred when the

buyer chooses the correct action, which can be written as:

$$\begin{aligned} c(I(v_b); \tau) &= \mathbb{E}_{\mathcal{R}, \mathcal{X}} [u_s^-(\sigma(\theta_r^{v_b}), x; \tau)] \\ &= \tau v_s + \tau(1 - 2v_s)(1 - v_s) \\ &\quad - \tau(1 - 2v_s)I(v_b)(v_s + \mathbf{1}_{\{I(v_b) \geq 0\}}), \end{aligned} \quad (9)$$

for which a full derivation is provided in Appendix C.3.

Mechanism Design. The objective of the mechanism is to maximize the seller's expected profit, i.e., the expected revenue from the transfer minus the expected cost, with the expectation taken over the type space of the buyer.

Proposition 4.3 (Optimal Mechanism). *The optimal mechanism maximizes the seller's profit subject to individual rationality and incentive compatibility, which can be expressed as the following mathematical program:*

$$\max_{I, t} \quad \mathbb{E}_{\mathcal{V}_b} [t(v_b) - c(I(v_b); \tau)] \quad (10a)$$

$$s.t. \quad I(v_b) \in \mathbb{R}_{[-1, 1]}, t(v_b) \in \mathbb{R}_{\geq 0} \quad (10b)$$

$$\delta(I(v_b), v_b) - t(v_b) \geq 0 \quad (10c)$$

$$\delta(I(v_b), v_b) - t(v_b) \geq \delta(I(b), v_b) - t(b) \quad (10d)$$

where (10a) is the expected profit, (10b) encodes the bounds on the allocations and transfers, and (10c) and (10d) encode individual rationality and incentive compatibility, respectively. Each constraint is defined for every $v_b \in \mathcal{V}_b$, as well as every $b \in \mathcal{B}$ for the incentive constraint (10d).

We note that, it is the inclusion of the cost described in (9) in the objective that generalizes existing works by allowing us to model the purchase of information from a competitor as opposed to a third party. In the following, we solve this problem for binary types to build intuition into the optimal menu, then generalize to continuous type spaces.

4.1. Binary Type Space

Suppose $\mathcal{V}_b = \{v_b^l, v_b^h\}$, with the high type v_b^h assigning a higher valuation to the fully informative communication rule than the low type, such that $\delta(0, v_b^h) > \delta(0, v_b^l)$. Without additional information, each type chooses the action that matches the state with the highest prior probability, meaning in this case the high type places lower prior probability on the state they deem most likely. In other words, the high type has lower precision (i.e., is less well informed), where we define the precision of a type as its distance from $1/2$, such that $|v_b^h - 1/2| < |v_b^l - 1/2|$. We denote the probability of the high type as $\phi = P(v_b^h)$.

Corollary 4.4. *From Proposition 4.3, with binary types the*




Figure 5. A congruent binary type distribution. Without additional information, both types will take the same action, that is, $\sigma(v_b^l) = \sigma(v_b^h) = 0$.

the mechanism design problem in (10) reduces to:

$$\max_{I^k, t^k} \phi(t^h - c^h) + (1 - \phi)(t^l - c^l) \quad (11a)$$

$$\text{s.t. } I^k \in \mathbb{R}_{[-1,1]}, t^k \in \mathbb{R}_{\geq 0} \quad (11b)$$

$$\delta(I^k, v_b^k) - t^k \geq 0 \quad (11c)$$

$$\delta(I^k, v_b^k) - t^k \geq \delta(I^{k'}, v_b^k) - t^{k'} \quad (11d)$$

where $I^k = I(v_b^k)$, $t^k = t(v_b^k)$ and $c^k = c(I^k)$ denote the allocation, transfer and externality cost for type $k \in \{l, h\}$, respectively, and the objective in (11a) denotes the expected profit. Constraints (11d) and (11c) are the incentive and individual rationality constraints, respectively. Constraints are defined for every $k \in \{l, h\}$, as well as every $k' \in \{l, h\}$ where $k \neq k'$ for the incentive constraint (11d).

We now distinguish between congruent and noncongruent types; two types are congruent if they would select the same action without additional information.

Congruent Types. Hereafter, we will assume that the high type selects $\sigma(v_b^h) = 0$ based on their prior information (i.e., $v_b^h > 1/2$). However, our results can be easily adapted to the opposite case. The types are congruent if $1/2 < v_b^h < v_b^l$, as illustrated in Figure 5. In this setting, by Proposition 4.1, so as to not reduce the information of the buyer relative to not participating, it is necessary that $I^l, I^h \geq 0$, meaning that the state $X = 0$ must be revealed with certainty. Moreover, as is typical in monopolistic screening problems, both the individual rationality constraint for v_b^l and the incentive compatibility constraint for v_b^h bind. Under these conditions, the problem in (11) simplifies to the following linear program:

$$\max_{I^l, I^h} \phi(t^h - c^h) + (1 - \phi)(t^l - c^l)$$

$$\text{s.t. } 0 \leq I^l \leq 1$$

$$0 \leq I^h \leq 1,$$

with $t^l = (1 - I^l)(1 - v_b^l)$ and $t^h = (1 - v_b^h)(I^l - I^h) + t^l$. This problem can be solved analytically, revealing that it is optimal for the seller to either allocate all information or none at all, such that both $I^l, I^h \in \{0, 1\}$. The optimal menu depends on the distribution of buyer types, as well as

the magnitude of competition encoded by τ and the seller's own beliefs v_s .

The linear program yields two decision boundaries, τ^l and τ^h , that determine the sign of the allocation as a function of competition intensity for the low and high type, given by

$$\tau^l = \frac{(1 - v_b^l) - \phi(1 - v_b^h)}{(1 - \phi)(v_s - 1)(2v_s - 1)},$$

and



$$\tau^h = \frac{1 - v_b^h}{(v_s - 1)(2v_s - 1)},$$

with $\tau^l < \tau^h$, which we plot in Figure 6 for $v_b^l = 5/6$ and $v_b^h = 2/3$.

If the seller's beliefs are opposite to the buyer's (i.e., $v_s \leq 1/2$), the solution is $I^l = \mathbf{1}_{\{\tau > \tau^l\}}$ and $I^h = \mathbf{1}_{\{\tau > \tau^h\}}$. The seller believes the buyer would choose the wrong action without additional information. Therefore, irrespective of ϕ , there is a point at which v_s is so far in the opposite direction of v_b , that if τ is sufficiently high, both $I^l = 1$ and $I^h = 1$. Here, $\tau > \tau^h$ and the expected cost of revealing the state deemed most likely by the seller is too much relative to the expected transfer such that no information is sold at all, even to v_b^h . Information is only sold to v_b^l when the increase in expected transfer outweighs the increase in cost such that $\tau < \tau^l$. This happens when the intensity of competition τ is sufficiently low (Figure 6a) or probability of the high type ϕ is sufficiently high (Figure 6c).

If $v_s > 1/2$, the seller reveals the state to v_b^h irrespective of τ as their beliefs are now in agreement. Ideally, the seller would minimize costs by obfuscating only $X = 0$ (i.e., set $I^h < 0$) to force the buyer to pick the *wrong* action at the expense of social welfare, however this would violate the obedience constraint as the buyer's gain would be negative by following this message. Information is only sold to v_b^l when $\tau > \tau^l$. As the probability of the high type increases, $\phi \rightarrow 1$, the area enclosed by the solid black line in Figure 6c shrinks, and no information is sold to the low type. This ensures incentive compatibility by preventing the high type from mimicking the low type.

Noncongruent Types. The two types are noncongruent if $v_b^l < 1/2 < v_b^h$, as illustrated in Figure 7. As the low type would instead select $\sigma(v_b^l) = 1$ without additional information, $I^l \leq 0$ to ensure individual rationality. Without competition, revenue is maximized by offering full information to the high type and partial information to the low type, where the quality of this partial information is prescribed as that which makes the incentive compatibility constraint of the high type bind. However, this does not hold generally, and in our case, depending on the intensity of competition


 Figure 6. Decision boundaries τ^h (dashed) and τ^l (solid) for congruent types.

 Figure 7. A noncongruent binary type distribution. Without additional information, types take different actions, that is, $\sigma(v_b^l) = 1$ and $\sigma(v_b^h) = 0$.

τ , the seller may offer less information to both types. Therefore, for noncongruent types, the problem in (11) reduces to the following linear program:

$$\begin{aligned} \max_{I^l, I^h} \quad & \phi(t^h - c^h) + (1 - \phi)(t^l - c^l) \\ \text{s.t.} \quad & \underline{I}^l \leq I^l \leq 0 \\ & 0 \leq I^h \leq 1, \end{aligned}$$

with $t^l = v_b^l(1 + I^l)$, $t^h = (1 - v_b^h)(1 - I^h)$ and $\underline{I}^l = (2v_b^h - 1)/(v_b^h - v_b^l) - 1$. This problem can also be solved analytically, revealing that $I^l = \underline{I}^l(1 - \mathbf{1}_{\{\tau > \tau^l\}})$ and $I^h = \mathbf{1}_{\{\tau > \tau^h\}}$ where


$$\tau^l = \frac{v_b^l}{v_s(2v_s - 1)},$$

and

$$\tau^h = \frac{1 - v_s^h}{(1 - v_s)(1 - 2v_s)},$$

which we plot in Figure 8 for $v_b^l = 1/6$ and $v_b^h = 2/3$.

If $v_s < 1/2$, there exists a point where v_s diverges sufficiently from v_b^h and $\tau > \tau^h$, so no information being sold to the high type. However, the seller may still offer partial information to the low type if it remains profitable without


 Figure 8. Decision boundaries τ^h (dashed) and τ^l (solid) for noncongruent types.

violating incentive constraints. The information offered to v_b^l depends on the relative precision (i.e., distance from $1/2$) of the types. As the low type's prior becomes less informative, approaching that of v_b^h , they receive more information due to their increased willingness to pay.

If $v_s > 1/2$, the high type receives full information, whilst the low type only if $\tau < \tau^l$. For both types, the boundary τ beyond which no information is provided decreases with precision, as precise types derive less value from additional information, so their willingness to pay no longer exceed the cost of providing the information.

4.2. Continuous Type Space

We now generalize to a type space that is continuous on the unit interval, such that the distribution $F(v_b)$ has support on $\mathcal{V}_b = \mathbb{R}_{[0,1]}$ with density $p(v_b)$. To obtain analytical insights, we begin with a key result from prior work, which extends

Myerson (1981)'s finding that incentive compatibility and individual rationality can be summarized by simpler, more compact constraints on the allocation $I(v_b)$ with transfers $t(v_b)$ pinned down by the utility of the lowest type. We then generalize this work by incorporating the seller's expected cost into the objective and derive closed-form results.

Proposition 4.5 (Bergemann et al., 2018). *Individual rationality and incentive compatibility hold as long as $I(v_b)$ is non-decreasing and $\int_{\mathcal{V}_b} I(v_b) dv_b = 0$.*

Proof. Provided in Appendix C (Sections C.1 and C.2). \square

Corollary 4.6. *From Proposition 4.5, the transfer function can be derived using Milgrom & Segal (2002)'s envelope theorem to be:*

$$t(v_b) = I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) - \int_0^{v_b} I(z) dz. \quad (12)$$

Following Myerson (1981), we can use integration by parts to eliminate the integral in (12) when computing $\mathbb{E}_{\mathcal{V}_b}[t(v_b)]$. Then, using the expression for the expected cost described in (9), we re-write the problem in (10) as follows.

Corollary 4.7. *From Proposition 4.5, the mechanism design problem in (10) reduces to:*

$$\max_I \mathbb{E}_{\mathcal{V}_b} [J(I, v_b)] \quad (13a)$$

$$s.t. \quad \int_{\mathcal{I}} I(v_b) dv_b = 0 \quad (13b)$$

$$\frac{d}{dv_b} I(v_b) \geq 0 \quad (13c)$$

where (13c) encodes the monotonicity constraint and

$$J(I, v_b) = I(v_b) \left(v_b + \frac{F(v_s)}{p(v_s)} + \tau v_s (1 - 2v_s) \right) + \mathbf{1}_{\{I(v_b) \geq 0\}} (\tau (1 - 2v_s)^2 - 1). \quad (14)$$

This formulation bears similarities to Myerson (1981)'s single item auction, where the problem is reduced to virtual surplus maximization subject to monotonic allocations. The virtual value function is defined as the partial derivative of the integrand in (14) with respect to I , given by

$$\begin{aligned} \pi(I, v_b) &= \frac{\partial}{\partial I} J(I, v_b) \\ &= \pi^+(v_b) + \mathbf{1}_{\{I(v_b) \leq 0\}} (\pi^-(v_b) - \pi^+(v_b)), \end{aligned}$$

where

$$\pi^-(v_b) = p(v_b) (\tau v_s (1 - 2v_s) + v_b) + F(v_b),$$

and

$$\pi^+(v_b) = p(v_b) (\tau (1 - v_s) (1 - 2v_s) + v_b - 1) + F(v_b),$$

where the negative π^- and positive π^+ parts represent the marginal benefit to the seller in terms of profit, albeit ignoring the constraint (13b), of increasing the allocation of each type from -1 to 0 and from 0 to 1 , respectively. Hereafter, we assume the virtual value function is non-decreasing in v_b , which places a regularity condition on the type distribution. This is a common assumption as it holds for several familiar distributions, including the uniform, exponential, and Gaussian distributions.

Irregularity. This regularity assumption excludes cases where most buyers are well-informed *ex-ante* (i.e., with types clustered near 0 or 1), resulting in a bimodal density that violates regularity, inducing virtual values that require *ironing*. Ironing, introduced by Myerson (1981), treats irregular distributions by smoothing non-monotonic intervals. Myerson's method transforms the virtual values π into a non-decreasing *ironed* version $\bar{\pi}$, ensuring that allocations are monotonic by design, as is the case for regular distributions. However, unlike in the classic model of monopolistic screening, the virtual values here are a function of I , rather than being constant in this decision variable, and so this method cannot be readily applied. One could adopt a similar method to Toikka (2011), which extends that of Myerson such that in our case

$$\bar{\pi}(I, v_b) = \frac{d}{dv_b} \text{conv} \left(\int_0^{v_b} \pi(z, I) dz \right) (v_b),$$

where

$$\begin{aligned} \int_0^{v_b} \pi(z, I) dz &= F(v_b) (\tau v_s (1 - 2v_s) \\ &\quad - \mathbf{1}_{\{I(v_b) \leq 0\}} (2\tau v_s (1 - 2v_s) + 1) + v_b), \end{aligned}$$


however this requires the virtual surplus to be weakly concave in I , which we shall show later does not always hold in our case.⁶ Therefore, we leave a thorough extension of our model to irregular type distributions to future work.

Solution. With the regularity condition in place, we deal with the equality constraint (13b) using standard Lagrange methods, such that the Lagrangian function is given by

$$L(I, \lambda) = \int_{\mathcal{V}_b} (J(I, z) - \lambda z) dz.$$

As $I(v_b)$ is non-decreasing in v_b , $J(I, v_b)$ is piecewise linear in I , and the integrals in (13) are linear functionals, strong duality holds thus $\lambda^* = \arg \min_{\lambda} \max_I L(I, \lambda)$ and hence we find I^* by maximizing $L(I, \lambda^*)$. By applying the results

⁶The operator $\text{conv}(\cdot)$ denotes the convex envelope. The convex envelope $\text{conv}(g)$ of a function $g : \mathbb{R}^n \mapsto \mathbb{R}$ is the largest convex underestimator of g whose epigraph is the convex hull of the epigraph of g , that is, $\text{epi}(\text{conv}(g)) = \text{conv}(\text{epi}(g))$.


 Figure 9. Virtual surplus with $v_b = 2/5$ for varying levels of competition.

of Myerson (1981), we can instead maximize virtual surplus pointwise, so the menu is optimal if there exists λ^* with

$$I^*(v_b) = \arg \max_{I \in \mathcal{I}} I(v_b) (\pi(I, v_b) - \lambda^*). \quad (15)$$

Proposition 4.8. *For regular type distributions, the solution to the mechanism design problem in (13) is such that the optimal menu comprises no partially informative communication rules, with $I^*(v_b) \in \{-1, 0, 1\}$.*

Proof. Provided in Appendix D. \square

Corollary 4.9. *Let $\tau' = (1 - 2v_s)^{-2}$, then provided $\tau < \tau'$, from Proposition 4.8, it holds that*

$$\lambda^* = \frac{\pi^-(1/2) + \pi^+(1/2)}{2},$$

and the primal solution to (15) is given by

$$I^*(v_b) = \begin{cases} -1 & \text{if } \pi^-(v_b) < \lambda^*, \\ 0 & \text{if } \pi^-(v_b) > \lambda^* > \pi^+(v_b), \\ 1 & \text{if } \pi^+(v_b) > \lambda^*, \end{cases}$$

yet, if $\tau \geq \tau'$, then $\lambda^* = \pi^+(1/2)$ and the primal solution reduces to

$$I^*(v_b) = \begin{cases} -1 & \text{if } v_b \leq 1/2, \\ 1 & \text{otherwise,} \end{cases}$$

meaning it is optimal for the seller to reveal no information at all.

To gain some insight into these results, suppose the buyer's types follow a uniform distribution. In this case, $F(v_b) = v_b$ and $p(v_b) = 1$ for every $v_b \in \mathbb{R}_{[0,1]}$, and the virtual values are given by

$$\pi^-(v_b) = 2v_b + \tau v_s(1 - 2v_s),$$

and

$$\pi^+(v_b) = 2v_b + \tau(1 - v_s)(1 - 2v_s) - 1.$$

Insights. In Figure 9, we plot the objective value without the integral constraint (i.e., the virtual surplus) for varying levels of competition, with $v_b = 2/5$. Irrespective of τ , the virtual surplus is piecewise linear in I . When $v_s = 1/2$, it is completely unaffected by τ . This is because the seller places equal probability on either state, so the expected externality is equal whether or not the buyer changes actions. Note that, this would not be the case if τ was a function of the state, as the expected cost could again be asymmetric.

If $v_s \leq 1/2$, both $\sigma(v_b) = 1$ and $\sigma(v_s) = 1$, so the buyer takes the same action as the seller, without additional information. In this case, when τ is low, the buyer can charge enough for it to be profitable to reveal the state. As competition intensifies, the seller wants to make the buyer choose the *wrong* action by obfuscating $X = 1$ and revealing only $X = 0$, so the offering changes to $I^*(v_b) = 1$. However, this would violate individual rationality so would not be possible with the integral constraint in place.

If $v_s \geq 1/2$, we have $\sigma(v_s) = 0$, so the buyer would choose the opposite action to the seller. In this case, as τ increases, the benefit of revealing the state becomes outweighed by the cost incurred when the buyer switches actions to the state they deem more likely, and eventually no information is sold. More generally, the threshold τ' is such that for every $\tau < \tau'$, the objective is concave, and the primal solution occurs at the kinks, $I^*(v_b) \in \{-1, 0, 1\}$, yet for $\tau \geq \tau'$, the objective is convex and the primal solution instead occurs at the endpoints, $I^*(v_b) \in \{-1, 1\}$. It is at this point when no transfer can be charged that justifies the externality cost of revealing the state.

In Figure 10, we plot the resulting primal solutions and the




Figure 10. Both negative (solid) and positive (dashed) virtual values, with $v_s = 4/5$, for varying levels of competition. The primal solution, $I^*(v_b)$, is indicated by the blue vertical lines, with only types within the enclosed interval offered the fully informative communication rule. The dual solution, λ^* , which identifies the threshold points, is highlighted in red.

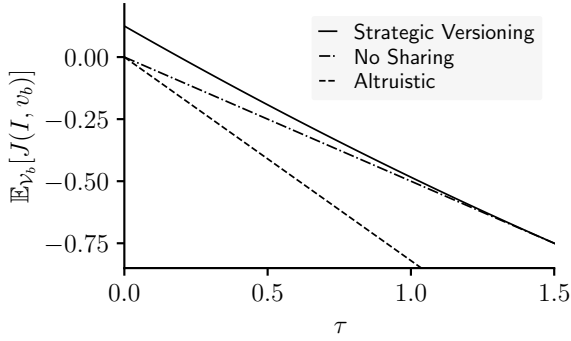


Figure 11. Expected profit from either strategic versioning (solid), offering full information to all types (dashed) or sharing no information at all (dashdot)

virtual values after re-introducing the integral constraint. To satisfy constraint (13b), the dual variable λ^* must assign two threshold types that separate those receiving full information, $I(v_b) = 0$, from those receiving no information at all, $I(v_b) \in \{-1, 1\}$. Partial information, $I(v_b) \in (-1, 1)$, is only offered to types where $\pi(I, v_b) = \lambda^*$. However, if the virtual value function is strictly increasing, which holds for all regular type distributions, then these thresholds are crossed exactly once, in other words, the virtual values have measure zero at these points. Thus, partial information is never offered and the optimal menu is a step function. For instance when $\tau = 0$, the fully informative communication rule is offered if $1/4 \leq v_b \leq 3/4$ at a price $t(v_b) = 1/4$. In this case, the results align with the setup of Bergemann et al. (2018), where the seller is not part of the game. The intuition is that it effectively corresponds to a setting with there

is no cost associated with the buyer's actions. Nevertheless, even in this case, Corollary 4.9 extends these prior results by providing an analytical solution for λ^* .

Since the objective is concave when $\tau < \tau'$, by inspection we see that $I(v_b) = 1$ is preferable to $I(v_b) = -1$ only if $\pi^+(v_b) - \lambda > \lambda - \pi^-(v_b)$, so the two threshold types are specified by the unique λ^* that satisfies this inequality with equality. Moreover, as the integral in constraint (13b) has uniform intervals, the set of types where $I(v_b) = -1$ is forced have the same measure as the set where $I(v_b) = 1$, which on the unit interval is measure $1/2$, so the dual variable is ultimately pinned down by the equality constraint $\pi^+(1/2) - \lambda^* = \lambda^* - \pi^-(1/2)$. As τ increases, the set of types offered the fully informative communication rule becomes narrower, raising the price of information as the seller extracts more revenue from fewer types to maximize revenue whilst minimizing the externality cost.

When $\tau = \tau'$, the virtual value functions collapse onto one another, and beyond this point the objective is convex. As the seller cannot mislead the buyer to taking the *wrong* action with sufficient probability without violating individual rationality, they can do no better offering no information. The primal solution therefore always occurs at the endpoints, so in order to satisfy the constraint (13b), exactly half of the types are offered $I(v_b) = -1$ and the other half $I(v_b) = 1$. Therefore, λ^* now only prescribes a single threshold type, where $I(v_b) = 1$ if $\pi^+(v_b) > \lambda$, else $I(v_b) = -1$, implying a threshold of $v_b = 1/2$ and dual solution $\lambda^* = \pi^+(1/2)$.

To round off our analyses, in Figure 11 we plot the expected profit obtained via (i) strategic versioning using our mechanism design approach; (ii) offering the fully informative communication rule to all types (i.e., that which would be obtained with altruistic information sharing); and (iii) shar-

ing no information at all. One can see that the seller is always better off sharing no information than revealing it all for free, due to the anti-coordination incentives in our setup. By strategically screening buyer types, our market mechanism provides incentives for information sharing since for low levels of competition, the profit earned is strictly positive. For moderate τ , selling information to some buyers is still better than not sharing at all. However, as competition intensifies, less and less information is sold, until eventually, none at all.

5. Conclusions

This paper examines a setting in which a monopolist sells supplemental information to a privately informed buyer. As with previous work, both the design and price of information are shaped by the buyer's prior beliefs. We extend this framework to account for the seller's own private information, the potential for the buyer and seller to be competitors in a downstream market, and the intensity of competition therein. Without product versioning, the seller is better off sharing no information than altruistically revealing the true state, and vice versa for the buyer—an outcome consistent with the efficient market hypothesis, as freely disclosing information would erode the seller's information rent. Hence, we show that with our mechanism design the seller can screen buyer types to maximize profit in a way that benefits both parties. That said, in fiercely competitive environments, the seller may still be better off not sharing as the externality cost outweighs the cost of doing so, since the transfers are ultimately pinned down by the demand for information and obedience constraints restrict the seller's ability to steer the buyer's actions at the expense of social welfare. In this case, regulatory measures or other market interventions may be required to incentivize sharing.

In this work, we characterized the profit-maximizing mechanism within a linear model with binary states and actions, hence much work remains to make this framework useful in practice. A natural next step is to extend our model to richer settings, with more than two firms, as well as multiple, or continuous, states and actions. Further, in many real-world markets, information induces nonlinear externalities that depend on all player's actions, leading to Bayesian Nash equilibria rather than dominant strategies. With multiple firms, one can model the complex network of externalities amongst buyer's and seller's alike, possibly represented by a weighted directed graph. Analyzing such environments may require computational approaches to characterize the optimal mechanism, leveraging recent advancements in learning-based methods for automated mechanism design.

On a broader note, whilst our results offer first pass insights into the sale of information to a competitor, our information design-based setup has its limitations. First, the seller is as-

sumed to be risk-neutral, yet in practice sellers would have different risk appetites, so incorporating risk measures (e.g., value at risk, expected shortfall, etc.) in the objective should be explored. Whilst economic literature often assumes perfect knowledge of the buyer's type distribution, this is rarely the case in practice, so future work could explore the impact of distributionally robust mechanism design or methods for learning the distribution on the design and price of information.

On the topic of learning, we focus only on static mechanisms, however in dynamic settings the seller could extract more surplus by correlating payments with realized states. Also in our setup, buyers are distinguished solely by differences in their beliefs. Yet, in real-world scenarios, they may exhibit heterogeneity along multiple dimensions, including their capacity for processing information, or their preference for timely access to information. Lastly, since the realized state is non-contractible, the seller cannot use scoring rules to price information, which raises the question of how much profit could be earned if this were the case.

A. Proof of Proposition 3.2

Suppose that the set of messages the seller can send to the buyer was $\mathcal{M} = \cup_{a_b} \mathcal{M}^{a_b}$, where \mathcal{M}^{a_b} is the subset of messages that incite type v_b to choose action $a_b \in \mathcal{A}_b$. Further, consider an alternative message space $\mathcal{R} = \cup_{a_b} \{r^{a_b}\}$, such that $|\mathcal{R}| = |\mathcal{A}_b|$, and each message r^{a_b} inciting type v_b to choose a different action.

In summary, for each action, set \mathcal{M} contains many messages that incite that action, whilst set \mathcal{R} contains only one.

Lemma A.1. *The distributions over \mathcal{M} and \mathcal{R} are such that $p(r^{a_b}|x; v_b)$ is a garbling of $p(m^{a_b}|x; v_b)$.*

Proof. If $p(r^{a_b}|x; v_b)$ is a garbling of $p(m^{a_b}|x; v_b)$, there must exist a function $g : \mathcal{R} \times \mathcal{M}^{a_b} \mapsto \mathbb{R}_{[0,1]}$ such that

$$p(r^{a_b}|x; v_b) = \sum_{m^{a_b} \in \mathcal{M}^{a_b}} g(r^{a_b}, m^{a_b}) p(m^{a_b}|x; v_b),$$

for every $a_b \in \mathcal{A}_b$, which in our case holds by design simply for $g(r^{a_b}, m^{a_b}) = 1$, for every $m^{a_b} \in \mathcal{M}^{a_b}$. \square

With Lemma A.1, message r^{a_b} conditioned on the state has equal probability mass to the sum of all those in \mathcal{M}^{a_b} , so the joint distribution of states and actions are the same, so v_b is indifferent between these two communication rules. In addition, garbling can be viewed as a way of adding noise to a communication rule, so is weakly less informative for all types. Hence, by Blackwell's theorem we get that

$$\mathbb{E}_{\mathcal{R}, \mathcal{X}} [u_b^+(\sigma(\theta_r^b), x)] \leq \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)],$$

for every $b \in \mathcal{B}$, which means that every other reported type weakly prefers $p(m|x; b)$ over $p(r|x; b)$, whilst the true type is indifferent (Blackwell, 1951; 1953).

In brief, considering only direct communication rules with $\mathcal{M} = \mathcal{A}_b$, the value will be unchanged for v_b , and the value will only be reduced for $b \neq v_b$, so incentive compatibility and individual rationality are preserved. \square

B. Proof of Proposition 4.1

We use the following result from convex analysis:

Lemma B.1 (Boyd, 2004). *Let $g : \mathbb{R}^n \mapsto \mathbb{R}$ be an affine function and let \mathcal{S} be a (nonempty) convex polytope. Then for every $z \in \text{int}(\mathcal{S})$ there exists a point z' on the boundary $\partial\mathcal{S}$ of \mathcal{S} with $p(z) = p(z')$.*

Proof. If g is constant on \mathcal{S} the claim is trivial, so assume that g is nonconstant. Then, for every $z \in \text{int}(\mathcal{S})$, the level set $\mathcal{L} = \{z' \in \mathbb{R}^n : g(z') = g(z)\}$ is the intersection of \mathcal{S} with an affine hyperplane. As \mathcal{S} is compact and convex, \mathcal{L} is nonempty, intersecting $\partial\mathcal{S}$ at some point z' . \square

Let $I_0 = P(m_0|X = 0)$ and $I_1 = P(m_1|X = 1)$ be a direct communication rule as in Table 1, with $I_0, I_1 \in \mathbb{R}_{[0,1]}$ and $I_0 + I_1 \geq 1$. Recall $I = I_0 - I_1$, so we can write the feasible region as $\mathcal{S} = \{(I_0, I) \in \mathbb{R}^2 : 0 \leq I_0 \leq 1, -1 \leq I \leq 1, I \geq 1 - 2I_0, I_0 + I \leq 1\}$ which defines a convex polytope. As the seller's objective is linear in these variables, by Lemma B.1, there exists a level set \mathcal{L} whose intersection with \mathcal{S} is nonempty. Therefore $\mathcal{L} \cup \mathcal{S}$ is a line segment with endpoints on $\partial\mathcal{S}$.

Let $\bar{\partial}\mathcal{S} = \{(I_0, I) \in \mathcal{S} : I_0 = h(I)\}$ be the ceiling of \mathcal{S} , with $h(I) = 1$ if $I \leq 0$, otherwise $h(I) = 1 - I$, as shown in Figure 3. We claim any line \mathcal{L} which intersects $\partial\mathcal{S}$ must intersect $\bar{\partial}\mathcal{S}$. If $\mathcal{L} \cup \bar{\partial}\mathcal{S} = \emptyset$, then $\mathcal{L} \cup \partial\mathcal{S}$ would lie entirely in the region where $I_0 < h(I)$. However, by the geometry of \mathcal{S} , any line traversing the boundary $\partial\mathcal{S}$ must eventually cross $h(I)$, thus $\mathcal{L} \cup \bar{\partial}\mathcal{S} \neq \emptyset$. Any feasible objective value has an equivalent solution on the ceiling, so communication rules can be defined by this line. \square

B.1. Proof of Corollary 4.2

If type v_b reports bid b , the expected value of the nonnegative utility term given message m is received can be written as follows:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)|m] &= \sum_{x \in \mathcal{X}} p(x|m, s_b; \theta_m^b) u_b^+(\sigma(\theta_m^b), x) \\ &= \theta_m^b \mathbf{1}_{\{\theta_m^b \geq 1/2\}} + (1 - \theta_m^b) \mathbf{1}_{\{\theta_m^b < 1/2\}} \\ &= (\theta_m^b \vee (1 - \theta_m^b)), \end{aligned}$$

and so by integrating over the message space, the expected nonnegative utility term is given by

$$\begin{aligned} \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] &= \mathbb{E}_{\mathcal{M}} [\mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)|m]] \\ &= \sum_{m \in \mathcal{M}} \left(v_b P(m|X = 0; b) \right. \\ &\quad \left. \vee (1 - v_b) P(m|X = 1; b) \right). \end{aligned}$$

Following Proposition 3.2, let $\mathcal{M} = \{m_0, m_1\}$, where m_0 and m_1 recommend actions $a_b = 0$ and $a_b = 1$. Let $I_0(b) = P(m_0|X = 0; b)$ and $I_1(b) = P(m_1|X = 1; b)$, then with inequalities (4) and (5), the above expression becomes

$$\begin{aligned} \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] &= v_b I_0(b) + (1 - v_b) I_1(b) \\ &= I_1(b) + v_b (I_0(b) - I_1(b)) \\ &= I_1(b) + v_b I(b), \end{aligned}$$

and following Proposition 4.1, we get

$$\begin{aligned} I_1(b) &= I_0(b) - (I_0(b) - I_1(b)) \\ &= I_0(b) - I(b) \\ &= 1 - I(b) \mathbf{1}_{\{I(b) \geq 0\}}, \end{aligned}$$

and so by substituting this into the previous expression, the gain of a communication rule is given by

$$\begin{aligned} \delta(I(b), v_b) &= \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_b^+(\sigma(\theta_m^b), x)] - \mathbb{E}_{\mathcal{X}} [u_b^+(\sigma(v_b), x)] \\ &= 1 - I(b) (\mathbf{1}_{\{I(b) \geq 0\}} - v_b) - \underbrace{(v_b \vee (1 - v_b))}_{\text{see (7)}}, \end{aligned}$$

for every $v_b, m \in \mathbb{R}_{[0,1]}$. \square

C. Proof of Proposition 4.5

In mechanism design, for a social choice function that is defined over a finite set of alternatives, monotonicity is, in general, necessary but not sufficient for implementability (Rochet, 1987). Identifying domains where monotonicity is sufficient remains an active research area (Ashlagi et al., 2010). However, Myerson (1981) proved sufficiency when a seller auctions a single item to buyers with private valuations and quasilinear utilities. We leverage similarities of this setup to ours, where I is the item allocation. Our proof is similar to that of Bergemann et al. (2018) as the demand for information is independent of the seller's cost. Nevertheless, we provide a detailed proof to be self-contained.

C.1. Necessity

To prove necessity, we start by assuming I is implementable, in which case truthfulness implies that

$$\begin{aligned} & \delta(I(v_b), v_b) - t(v_b) \\ & \geq \delta(I(b), v_b) - t(b) \\ & = \delta(I(b), b) - t(b) + \delta(I(b), v_b) - \delta(I(b), b), \end{aligned}$$

for every $b \in \mathcal{B}$. If we now define $\Delta(z) = \delta(I(z), z) - t(z)$ as the rent of type $z \in \mathcal{V}_b$ for a truthful report, it follows from (8) that we can re-write this inequality as

$$\begin{aligned} & \Delta(v_b) - \Delta(b) \\ & = \delta(I(v_b), v_b) - t(v_b) - \delta(I(b), b) + t(b) \\ & \geq \delta(I(b), v_b) - \delta(I(b), b) \\ & = (v_b - b)I(b) - (v_b \vee (1 - v_b)) + (b \vee (1 - b)). \end{aligned}$$

Without loss of generality we can assume that $v_b \geq b$, as we can simply swap terms otherwise. Then, for $v_b < 1/2$, this inequality can be re-written as

$$I(v_b) + 1 \geq \frac{\Delta(v_b) - \Delta(b)}{v_b - b} \geq I(b) + 1$$

and so $I(v_b) \geq I(b)$ if $v_b \geq b$, hence monotonicity is a necessary condition, which can be shown to hold for $v_b > 1/2$ in a similar fashion.

C.2. Sufficiency

Next, to prove sufficiency we need to consider the fact that Δ is differentiable with respect to v_b on $[0, 1/2)$ and $(1/2, 1]$, so we examine these two intervals separately. Within these intervals, the value is continuous in v_b , so by the envelope theorem (Milgrom & Segal, 2002), incentive compatibility imposes, for $v_b \leq 1/2$

$$\begin{aligned} \Delta(1/2) &= \Delta(0) + \int_0^{1/2} \frac{d}{dv_b} \delta(I(v_b), v_b) dv_b \\ &= \Delta(0) + \int_0^{1/2} (I(v_b) + 1) dv_b, \end{aligned}$$

and similarly for $v_b \geq 1/2$

$$\begin{aligned} \Delta(1/2) &= \Delta(1) - \int_{1/2}^1 \frac{d}{dv_b} \delta(I(v_b), v_b) dv_b \\ &= \Delta(1) - \int_{1/2}^1 (I(v_b) - 1) dv_b, \end{aligned}$$

therefore, as information has no value for types $v_b \in \{0, 1\}$, such that $\Delta(0) = \Delta(1) = 0$, the following must hold

$$\begin{aligned} v_b|_0^{1/2} + \int_0^{1/2} I(v_b) dv_b &= v_b|_{1/2}^1 - \int_{1/2}^1 I(v_b) dv_b \\ \implies \int_0^1 I(v_b) dv_b &= 0, \end{aligned}$$

thus constraint (13b) is needed. As $\Delta(v_b) = \delta(I(v_b), v_b) - t(v_b)$, we can construct the following transfers from the envelope representation if $v_b \leq 1/2$,

$$\begin{aligned} t(v_b) &= \delta(I(v_b), v_b) - \int_0^{v_b} (I(z) + 1) dz \\ &= I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) - \int_0^{v_b} I(z) dz, \end{aligned}$$

and similarly if $v_b \geq 1/2$,

$$\begin{aligned} t(v_b) &= \delta(I(v_b), v_b) + \int_{v_b}^1 (I(z) - 1) dz \\ &= I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) + \int_{v_b}^1 I(z) dz \\ &= I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) - \int_0^{v_b} I(z) dz, \end{aligned}$$

where the last line is due to (13b). Now we have an expression the transfers for every $v_b \in \mathcal{V}_b$, so we write the rent of type v_b from reporting b as follows:

$$\begin{aligned} & \delta(I(b), v_b) - t(b) \\ &= 1 - I(b) (\mathbf{1}_{\{I(b) \geq 0\}} - v_b) - (v_b \vee (1 - v_b)) \\ & \quad - I(b) (b - \mathbf{1}_{\{I(b) \geq 0\}}) + \int_0^b I(z) dz \\ &= 1 - (v_b \vee (1 - v_b)) + (v_b - b)I(b) + \int_0^b I(z) dz. \end{aligned}$$

By inspection, if $b < v_b$ then the term $(v_b - b)$ is positive and the rent increases as $b \rightarrow v_b$ as I is monotone. On the other hand, if $b > v_b$, the term $(v_b - b)$ is negative so we want to decrease I , which happens as $b \rightarrow v_b$. The rent is thus maximized at $b = v_b$, hence the incentive constraints are satisfied. Individual rationality is also satisfied given the resulting rent is nonnegative for all $v_b \in \mathcal{V}_b$, hence we have proved sufficiency.

C.3. Proof of Corollary 4.7

Lastly, we show that the seller's objective can be reduced to that in (14). Stated in Section 4, for a given I , the seller's cost is the expected externality when the buyer chooses the correct action, which for a given state x , is:

$$\begin{aligned} \mathbb{E}_{\mathcal{X}} [u_s^-(\sigma(\theta_m^b), x; \tau) | m] &= \sum_{x \in \mathcal{X}} p(x | m, s_s; v_s) u_s^-(\sigma(\theta_m^b), x; \tau) \\ &= \tau v_s \mathbf{1}_{\{\theta_m^b \geq 1/2\}} + \tau (1 - v_s) \mathbf{1}_{\{\theta_m^b < 1/2\}}. \end{aligned}$$

Following Proposition 3.2, we know that $\theta_m^b \geq 1/2$ for m_0 and $\theta_m^b < 1/2$ for m_1 , so by integrating over the message

space, the expected cost given a communication rule is

$$\begin{aligned} c(I(b); \tau) &= \mathbb{E}_{\mathcal{M}, \mathcal{X}} [u_s^-(\sigma(\theta_m^b), x; \tau)] \\ &= \mathbb{E}_{\mathcal{M}} [\mathbb{E}_{\mathcal{X}} [u_s^-(\sigma(\theta_m^b), x; \tau) | m]] \\ &= \tau v_s P(m_0; v_s) + \tau(1 - v_s) P(m_1; v_s) \end{aligned}$$

and given we know that

$$\begin{aligned} P(m_1; v_s) &= v_s I_1(b) + (1 - v_s) I_1(b) \\ &= v_s(1 - I_0(b)) + (1 - v_s) I_1(b) \\ &= v_s + I_1(b) - v_s I(v_s) - 2v_s I_1(b) \\ &= v_s - v_s I(b) + (1 - 2v_s) I_1(b) \\ &= 1 - v_s - v_s I(b) - (1 - 2v_s) (I(b) \vee 0) \end{aligned}$$

the expected cost reduces to

$$\begin{aligned} c(I(b); \tau) &= \tau v_s + \tau(1 - 2v_s)(1 - v_s) \\ &\quad - \tau(1 - 2v_s) I(b) (v_s + \mathbf{1}_{\{I(b) \geq 0\}}), \end{aligned}$$

such that considering direct mechanisms, the sellers minimizes $\mathbb{E}_{\mathcal{V}_b} [c(I(v_b); \tau)]$. Now we turn our attention to the transfers, with the expected transfer given by

$$\begin{aligned} \mathbb{E}_{\mathcal{V}_b} [t(v_b)] &= \int_{\mathcal{V}_b} t(v_b) dF(v_b) \\ &= \int_{\mathcal{V}_b} \left(I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) \right. \\ &\quad \left. - \int_0^{v_b} I(v_b) dv_b \right) dF(v_b), \end{aligned}$$

hence if we let $\alpha = \int_0^{v_b} I(v_b) dv_b$ and $d\beta = p(v_b) dv_b$, then the right-most expression reduces to

$$\begin{aligned} &\int_{\mathcal{V}_b} \left(\int_0^{v_b} I(v_b) dv_b \right) p(v_b) dv_b \\ &= \int_{\mathcal{V}_b} \alpha d\beta \\ &= \alpha \beta \Big|_0^1 - \int_{\mathcal{V}_b} \beta d\alpha \\ &= F(v_b) \int_0^{v_b} I(z) dz \Big|_0^1 - \int_{\mathcal{V}_b} F(v_b) I(v_b) dv_b \\ &= \int_{\mathcal{V}_b} I(v_b) dv_b - \int_{\mathcal{V}_b} F(v_b) I(v_b) dv_b \\ &= \int_{\mathcal{V}_b} I(v_b) (1 - F(v_b)) dv_b, \end{aligned}$$

which if we substitute back into the original expression gives

$$\begin{aligned} \mathbb{E}_{\mathcal{V}_b} [t(v_b)] &= \int_{\mathcal{V}_b} \left(I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) \right) dF(v_b) \\ &\quad - \int_{\mathcal{V}_b} I(v_b) (1 - F(v_b)) dv_b \\ &= \int_{\mathcal{V}_b} \left(I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) \right) dF(v_b) \\ &\quad - \int_{\mathcal{V}_b} I(v_b) \frac{1 - F(v_b)}{p(v_b)} dF(v_b) \\ &= \int_{\mathcal{V}_b} \left(I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) - I(v_b) \frac{1 - F(v_b)}{p(v_b)} \right) dF(v_b) \\ &= \int_{\mathcal{V}_b} \left(I(v_b) (v_b - \mathbf{1}_{\{I(v_b) \geq 0\}}) + I(v_b) \frac{F(v_b)}{p(v_b)} \right) dv_b \\ &\quad - \underbrace{\int_{\mathcal{V}_b} I(v_b) dv_b}_{=0 \text{ by (13b)}} \\ &= \int_{\mathcal{V}_b} \left(I(v_b) \left(v_b - \mathbf{1}_{\{I(v_b) \geq 0\}} + \frac{F(v_b)}{p(v_b)} \right) \right) dF(v_b). \end{aligned}$$

Finally, the expected profit is equal to the expected transfer minus the expected cost, so the objective function in (13) is given by

$$\mathbb{E}_{\mathcal{V}_b} [t(v_b) - c(I(v_b); \tau)] = \int_{\mathcal{V}_b} J(I, v_b) dF(v_b),$$

where

$$\begin{aligned} J(I, v_b) &= I(v_b) \left(v_b + \frac{F(v_b)}{p(v_b)} + \tau v_s (1 - 2v_s) \right. \\ &\quad \left. + \mathbf{1}_{\{I(v_b) \geq 0\}} (\tau(1 - 2v_s)^2 - 1) \right). \quad \square \end{aligned}$$

D. Proof of Proposition 4.8

By inspection, the solution to the Lagrangian maximization in (15) is given by

$$I^*(v_b) = \begin{cases} -1 & \text{if } \pi^-(v_b) < \lambda^*, \\ z \in \mathbb{R}_{[-1,0]} & \text{if } \pi^-(v_b) = \lambda^*, \\ 0 & \text{if } \pi^-(v_b) > \lambda^* > \pi^+(v_b), \\ z \in \mathbb{R}_{[0,1]} & \text{if } \pi^+(v_b) = \lambda^*, \\ 1 & \text{if } \pi^+(v_b) > \lambda^*, \end{cases}$$

such that partial information is only offered when the virtual values coincide with the dual variable. With the regularity assumption, the virtual values are strictly increasing, so they have measure zero at these points. Thus, partial information is never offered and the optimal menu is a step function, $I^*(v_b) \in \{-1, 0, 1\}$.

References

- Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256, 2022.
- Admati, A. R. and Pfleiderer, P. A monopolistic market for information. *Journal of Economic Theory*, 39(2): 400–438, 1986.
- Admati, A. R. and Pfleiderer, P. Selling and trading on information in financial markets. *The American Economic Review*, 78(2):96–103, 1988.
- Admati, A. R. and Pfleiderer, P. Direct and indirect sale of information. *Econometrica: Journal of the Econometric Society*, pp. 901–928, 1990.
- Agarwal, A., Dahleh, M., and Sarkar, T. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 701–726, 2019.
- Agarwal, A., Dahleh, M., Horel, T., and Rui, M. Towards data auctions with externalities. *Games and Economic Behavior*, 148:323–356, 2024.
- Ashlagi, I., Braverman, M., Hassidim, A., and Monderer, D. Monotonicity and implementability. *Econometrica*, 78(5):1749–1772, 2010.
- Aumann, R. J. Correlated equilibrium as an expression of bayesian rationality. *Econometrica: Journal of the Econometric Society*, pp. 1–18, 1987.
- Babaioff, M., Kleinberg, R., and Paes Leme, R. Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 92–109, 2012.
- Bergemann, D. and Bonatti, A. Markets for information: An introduction. *Annual Review of Economics*, 11:85–107, 2019.
- Bergemann, D. and Morris, S. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.
- Bergemann, D. and Morris, S. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- Bergemann, D., Bonatti, A., and Smolin, A. The design and price of information. *American economic review*, 108(1): 1–48, 2018.
- Bergemann, D., Cai, Y., Velezgas, G., and Zhao, M. Is selling complete information (approximately) optimal? In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 608–663, 2022.
- Bimpikis, K., Crapis, D., and Tahbaz-Salehi, A. Information sale and competition. *Management Science*, 65(6):2646–2664, 2019.
- Blackwell, D. Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 2, pp. 93–103. University of California Press, 1951.
- Blackwell, D. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pp. 265–272, 1953.
- Bonatti, A., Dahleh, M., Horel, T., and Nouripour, A. Selling information in competitive environments. *Journal of Economic Theory*, 216:105779, 2024.
- Boyd, S. Convex optimization. *Cambridge UP*, 2004.
- Cai, Y. and Velezgas, G. How to sell information optimally: An algorithmic study. *arXiv preprint arXiv:2011.14570*, 2020.
- Castiglioni, M., Bacchiocchi, F., Marchesi, A., Romano, G., and Gatti, N. Selling information while being an interested party. *arXiv preprint arXiv:2301.13790*, 2023.
- Clarke, R. N. Collusion and the incentives for information sharing. *The Bell Journal of Economics*, pp. 383–394, 1983.
- Crawford, V. P. and Sobel, J. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pp. 1431–1451, 1982.
- Cr  mer, J. and McLean, R. P. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pp. 1247–1257, 1988.
- Dahleh, M. A., Tahbaz-Salehi, A., Tsitsiklis, J. N., and Zoumpoulis, S. I. Coordination with local information. *Operations Research*, 64(3):622–637, 2016.
- European Commission. Digital markets act: Commission welcomes political agreement on rules ensuring fair and open digital markets, 2021. URL https://ec.europa.eu/commission/presscorner/detail/en/ip_21_3143. [Accessed: 2025-02-14].
- Falconer, T., Kazempour, J., and Pinson, P. Bayesian regression markets. *Journal of Machine Learning Research*, 25(180):1–38, 2024. URL <http://jmlr.org/papers/v25/23-1385.html>.

- Falconer, T., Kazempour, J., and Pinson, P. Towards replication-robust data markets, 2025. URL <https://arxiv.org/abs/2310.06000>.
- Fallah, A., Jordan, M. I., Makhdoumi, A., and Malekian, A. The limits of price discrimination under privacy constraints. *arXiv preprint arXiv:2402.08223*, 2024a.
- Fallah, A., Jordan, M. I., Makhdoumi, A., and Malekian, A. On three-layer data markets. *arXiv preprint arXiv:2402.09697*, 2024b.
- Fallah, A., Makhdoumi, A., Malekian, A., and Ozdaglar, A. Optimal and differentially private data acquisition: Central and local mechanisms. *Operations Research*, 72(3):1105–1123, 2024c.
- Feng, Z., Chen, J., and Zhu, Y. Uncovering value of correlated data: trading data based on iterative combinatorial auction. In *2021 IEEE 18th International Conference on Mobile Ad Hoc and Smart Systems (MASS)*, pp. 260–268. IEEE, 2021.
- Floridi, L. What is the philosophy of information? *Metaphilosophy*, 33(1-2):123–145, 2002.
- Floridi, L. Philosophical conceptions of information. In *Formal theories of information: From Shannon to semantic information theory and general concepts of information*, pp. 13–53. Springer, 2009.
- Gal-Or, E. Information sharing in oligopoly. *Econometrica: Journal of the Econometric Society*, pp. 329–343, 1985.
- Gradwohl, R. and Tennenholtz, M. Selling data to a competitor. *arXiv preprint arXiv:2302.00285*, 2023.
- Hellwig, C. Public information, private information, and the multiplicity of equilibria in coordination games. *Journal of Economic Theory*, 107(2):191–222, 2002.
- Hotelling, H. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929. ISSN 00130133, 14680297. URL <http://www.jstor.org/stable/2224214>.
- Kakhbod, A., Ozdaglar, A., and Schneider, I. Selling wind. *The Energy Journal*, 42(1):1–38, 2021.
- Kamenica, E. Bayesian persuasion and information design. *Annual Review of Economics*, 11:249–272, 2019.
- Kamenica, E. and Gentzkow, M. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Koutsopoulos, I., Gionis, A., and Halkidi, M. Auctioning data for learning. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*, pp. 706–713. IEEE, 2015.
- Mathevet, L., Perego, J., and Taneva, I. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020.
- Milgrom, P. and Segal, I. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- Milgrom, P. R. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pp. 380–391, 1981.
- Murphy, A. H. What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and forecasting*, 8(2):281–293, 1993.
- Myerson, R. B. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- Myerson, R. B. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pp. 323–358, 1986.
- Myerson, R. B. and Satterthwaite, M. A. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.
- Novshek, W. and Sonnenschein, H. Fulfilled expectations cournot duopoly with information acquisition and release. *The Bell Journal of Economics*, pp. 214–218, 1982.
- Pinson, P., Han, L., and Kazempour, J. Regression markets and application to energy forecasting. *Top*, 30(3):533–573, 2022.
- Raith, M. A general model of information sharing in oligopoly. *Journal of economic theory*, 71(1):260–288, 1996.
- Rasouli, M. and Jordan, M. I. Data sharing markets. *arXiv preprint arXiv:2107.08630*, 2021.
- Ravindranath, S. S., Jiang, Y., and Parkes, D. C. Data market design through deep learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Rochet, J.-C. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of mathematical Economics*, 16(2):191–200, 1987.
- Rodríguez Olivera, R. Strategic incentives and the optimal sale of information. *American Economic Journal: Microeconomics*, 16(2):296–353, 2024.
- Shapiro, C. Exchange of cost information in oligopoly. *The review of economic studies*, 53(3):433–446, 1986.
- Spence, M. Job market signaling. In *Uncertainty in economics*, pp. 281–306. Elsevier, 1978.

Taneva, I. Information design. *American Economic Journal: Microeconomics*, 11(4):151–185, 2019.

The Economist. The world’s most valuable resource is no longer oil, but data. *The Economist*, 2017. URL <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data>. [Accessed: 2024-12-25].

Toikka, J. Ironing without control. *Journal of Economic Theory*, 146(6):2510–2526, 2011.

Vives, X. Duopoly information equilibrium: Cournot and bertrand. *Journal of economic theory*, 34(1):71–94, 1984.

Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., and Gao, Y. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.

Ziv, A. Information sharing in oligopoly: The truth-telling problem. *The RAND Journal of Economics*, pp. 455–465, 1993.