Optimal Pricing of Information

Shuze Liu,[1] **Weiran Shen,**[2] **Haifeng Xu** [1*]

[1] University of Virginia
[2] Renmin University of China
sl5nw@virginia.edu, shenweiran@ruc.edu.cn, hx4ad@virginia.edu,

Abstract

A decision maker is deciding between an active action (e.g., purchase a house, invest certain stock) and a passive action. The payoff of the active action depends on the buyer's private type and also an unknown *state of nature*. An information seller can design experiments to reveal information about the realized state to the decision maker, and would like to maximize profit from selling such information. We fully characterize, in closed-form, the revenue-optimal information selling mechanism for the seller.

After eliciting the buyer's type, the optimal mechanism charges the buyer an upfront payment and then simply reveals whether the realized state passed a certain threshold or not. The optimal mechanism features both price discrimination and information discrimination. The special buyer type who is a priori indifferent between the active and passive action benefits the most from participating the mechanism.

Introduction

In numerous situations, a decision maker wishes to take an active move but is uncertain about its outcome and payoff. Such active moves range from financial decisions of investing a stock or startup to daily-life decisions of purchasing a house or a used car, from macro-level enterprise decisions of developing a new product to micro-level decisions of approving a loan applicant or displaying online ads to a particular Internet user. In all these situations, the decision maker's payoff for the active move relies on uncertain information regarding, e.g., potential of the invested company, quality of the house, popularity of the new product, credit of the loan applicant, etc. Certainly, the decision maker typically also has a passive backup option of not making the move, in which case he obtains a safe utility without any risk. To decide between the *active* and the *passive* action, the decision maker can turn to an information seller who can access more accurate information about the uncertainties and thus help to better estimate the payoff for his action. Given the value of the seller's information to the decision maker, the seller can make a profit from how much the information helped to improve utilities of the decision maker, i.e., the information buyer.

This paper studies how a monopolistic information *seller* (she) can design an optimal pricing mechanism to sell her information to an information *buyer* (he). The buyer (a decision maker) needs to take one of two actions. The *active* action results in a payoff $v(q,t)$ where t captures the buyer's private *type* and the *state of nature* q summarizes the payoff-relevant uncertainty unknown to the buyer. The *passive* action for the buyer always results in the same utility, normalized to 0, regardless of q, t. Both q and t are random variables drawn independently from publicly known distributions. That is, the type t captures the buyer's private preference and is assumed to be irrelevant to the informational variable q.[1] The seller can design experiments to reveal partial information about state q, and would like to design an optimal mechanism to sell her information to a buyer randomly drawn from the type distribution.

The problem of selling information turns out to differ significantly from the classic pricing problem for selling goods. First, when selling (physical or digital) goods, the seller's allocation rule can be described by a probability of giving out the goods and a risk-neutral buyer's utility is linear in the allocation variable. However, when revealing information to a buyer through experiments, the design variable of an experiment for each buyer type is high-dimensional or can even be a functional when the state is a continuum. Moreover, the buyer's utility is generally non-linear in the variables that describe an experiment (Bergemann and Morris 2019). Second, in selling goods, any individually rational buyer would participate as long as their expected utility is at least 0. However, in our setup of selling information, the buyer may already have positive utility from his active action even without participating in the mechanism. An individually rational buyer would participate in the mechanism only when his utility will become even higher. These differences make the seller's optimization task more challenging. This will be evident later in our characterization of the optimal mechanism, which turns out to be significantly different from, and ar-

*We thank EC'2021 reviewers and audiences for helpful comments. Haifeng Xu would like to thank Dirk Bergmann, Margaret Meyer, Yining Guo, and seminar participants of the 2021 Workshop on Strategic Communication and Learning for inspiring suggestions.

[1]This independence assumption is relaxed in the full arXiv version. https://arxiv.org/abs/2102.13289

guably more intricate than, the optimal pricing mechanism for selling goods by (Myerson 1981).

Main Result

We consider the above information selling problem and characterize in closed-form the revenue-optimal mechanism, among all *sequential mechanisms* that includes all possible ways through which the seller may sequentially reveal information and ask for payments. To simplify the exposition, we assume that the buyer's value function is *linear* and *monotone non-decreasing* in t, i.e., $v(q,t) = \alpha(q)[t + \beta(q)]$ for some $\alpha(q) \geq 0$ and $\beta(q)$. In the full arXiv version[2], we discuss how our analysis and results can be generalized to any convex and monotone (in t) value functions $v(q,t)$.

Assuming $v(q,t) = \alpha(q)[t + \beta(q)]$, we show that there always exists an optimal mechanism of a simple format — a multi-entry menu where each entry containing a threshold experiment and a payment for each buyer type. In this optimal mechanism, the buyer is incentivized to report his true type t first.[3] The seller then charges the buyer p_t and, afterwards, designs an experiment to reveal whether the realized state q satisfies $\beta(q) \geq \theta_t$ or not for some carefully chosen threshold θ_t. We thus call the mechanism a *threshold mechanism*. The thresholds and payments generally vary for different buyer types, and are carefully designed to accommodate the amount of risk each buyer type can tolerate. That is, the optimal mechanism features both price discrimination and information discrimination. We fully characterize the threshold and payment in the optimal mechanism. Depending on the setting, the negative of the threshold (i.e., $-\theta_t$) turns out to equal either the (*lower*) virtual value of type t as defined by (Myerson 1981), or its variant which we coin the *upper* virtual value, or a novel convex combination of both coined the *mixed* virtual value.

The above optimal mechanism exhibits multiple interesting properties. First, the optimal mechanism turns out to only need to price the experiment with one round information revelation, even though the seller in our model is allowed to price experiment outcomes (i.e., signals) and use multiple rounds of information revelation. This is due to the independence of the informational variable q and buyer type t, which makes an upfront payment and an "aggregrated" experiment without loss of generality. Second, the special buyer type $\bar{t}$ who is a-priori indifferent between active and passive action has the largest surplus from participating the mechanism. This is aligned with our intuition that this buyer type should benefit the most from additional information since the two actions appear indistinguishable to him a-priori. Moreover, we show that the buyer surplus as a function of his type t is increasing and convex when $t \leq \bar{t}$ but immediately transitions to be decreasing and convex when $t \geq \bar{t}$. However, the buyer payment may be increasing or decreasing in t, depending on the setting. Third, information discrimination turns out to be crucial for revenue. We show that if information discrimination is not allowed, i.e., suppose the same experiment must be used for all buyer types, then the best the seller can do in this case is to reveal full information and charge the Myerson's reserve price. We demonstrate via an example that the revenue in this case may be arbitrarily worse than the optimal. However, under the monotone hazard rate assumption of the buyer type distribution, we show that the optimal single-entry menu can always guarantee at least $1/e(\approx 0.368)$ fraction of the optimal revenue.

Related Works

The most related literature to our work is the recent study by (Bergemann, Bonatti, and Smolin 2018), who also consider selling information to a decision maker. In their model, the state of nature affects the payoff of every action. They characterize the optimal mechanism for the special cases with binary states and actions or with binary buyer types, whereas only partial properties about the optimal mechanism can be derived for the general case. In contrast, in our setup the state only affects the payoff of the buyer's active action. This restriction allows us to characterize the closed-form solution of the optimal mechanism with many (even continuous) states and buyer types, and for general buyer payoff functions. Moreover, our design space of mechanisms allows multiple rounds of information revelation and also allows contracting the experiment outcomes (i.e., realized signals), though it turns out that the optimal mechanism only needs to price one-round experiments.[4] While (Bergemann, Bonatti, and Smolin 2018) also restrict their design space to mechanisms that only price one-round experiments, they pointed out that this restriction does lose generality in their general setup. That is, the seller may derive strictly more revenue by using multi-rounds of experiments or by contracting the experiment outcomes.

(Eső and Szentes 2007b) studied the pricing of advice in a principal-agent model motivated by consulting. The principal as a consultant in their model can contract the agent's actions. With such strong bargaining power, their main result shows that even the principal observes completely irrelevant information about the agent's payoffs, the principal can still obtain revenue that is as high as in the situation where she fully observes the agents' payoffs. However, different from consulting service, our model of information selling assumes that only information itself (i.e., the experiment or the experiment outcomes) is contractible and the buyer's actions are not contractible. Therefore, the main result of (Eső and Szentes 2007b) clearly does not hold in our model — if the seller's information is irrelevant to the buyer's payoffs in our model, she will certainly get zero revenue. Interestingly, the format of our optimal mechanism turns out to bear somewhat similar structure to the optimal contract of (Eső and Szentes 2007b), however our results are derived through different techniques and apply to much more general buyer value functions, whereas (Eső and Szentes 2007b)

[2]The full arXiv version: https://arxiv.org/abs/2102.13289

[3]Equivalently, it is the best interest for each buyer type to choose the particular menu intended for him. That is, the mechanism is incentive compatible.

[4]This is first observed by (Babaioff, Kleinberg, and Paes Leme 2012) yet we will provide a formal argument later for completeness.

restrict to simpler agent utility functions (i.e., the sum of the agent type and the state) and only log-concave agent type distributions. (Hörner and Skrzypacz 2016) study the problem where a firm faces a decision on whether to hire an agent, who has a binary private type, i.e., competent or not. The firm and agent can interact for many rounds by making money transfer and taking test to elicit information about the agent's type. They analyze the equilibrium when the number of rounds of interactions grows large. Both the model and the nature of their results are different from us.

There has also been recent interest of algorithmic studies that formulate optimization programs to compute the optimal mechanism for selling information to a decision maker. (Babaioff, Kleinberg, and Paes Leme 2012) prove revelation principle types of results and characterize the format of the optimal mechanism, depending on whether the state and buyer type are correlated or not; they then develop optimization programs to compute the optimal mechanism. The efficiency for solving these programs were later improved by (Chen, Xu, and Zheng 2020).

Model and Problem Formulation

The Setup

We study the following optimal information pricing problem between an information *seller* (she) and an information *buyer* (he). The the buyer is a decision maker who faces one of two actions: a *passive* action 0 and an *active* action 1. The the buyer obtains an uncertain payoff $v(q,t)$ for the active action 1 where $q \in Q$ is a random *state of nature* unknown to the buyer and $t \in T$ is the buyer's private type. Both T, Q are measurable sets. The buyer's utility for the passive action 0 is always 0, irrespective of his type and the state of nature. In other words, the passive action is a backup option for the buyer. For example, if the the buyer is a potential purchaser of some goods (e.g., a house or a used car) with uncertain quality, the passive action 0 corresponds to "not purchase" in which case the buyer has no gain or loss, whereas the active action 1 corresponds to "purchase" under which the the buyer's utility depends on the quality q of the goods as well as how much he values the goods (captured by his private type t).

Both t and q are random variables that are independently distributed according to the cumulative distribution functions (CDF) $F(t)$ and $G(q)$, respectively. We assume throughout the paper that that both $F(t)$ and $G(q)$ are continuously differentiable, with corresponding probability density functions (PDF) $f(t)$ and $g(q)$. Both $F(t)$ and $G(q)$ are public knowledge. However, the realized q can only be observed by the information seller. We study the the seller's problem of designing a *revenue*-maximizing pricing mechanism to sell her private observation of q to the the buyer. Notably, the buyer's private type t is only known to himself — had the the seller known the buyer type t, the seller's optimal pricing mechanism is simply to reveal full information and then charge the buyer the *value of (full) information* (Bergemann, Bonatti, and Smolin 2018): $\int_{q\in Q} \max\{0, v(q,t)\}g(q)\mathrm{d}q - \max\{0, \int_{q\in Q} v(q,t)g(q)\mathrm{d}q\}$.

Throughout we assume that the buyer payoff function $v(q,t)$ is monotone non-decreasing in his type t for any $q \in Q$. For expositional simplicity, we will assume $v(q,t)$ is linear in t, i.e., there exist real-valued functions $\alpha(q) \geq 0$ and $\beta(q)$ such that $v(q,t) = \alpha(q)(t + \beta(q))$.

In our arXiv version, we show how our results and analysis easily generalize to any convex (in t) function $v(q,t)$. Linearity also implies that the buyer's type $t \in \mathbf{R}$ is a real value, which we assume is supported on a closed interval $T = [t_1, t_2]$ [5]. However, the state q is allowed to be supported on a general measurable set Q and does *not* need to be a real value. Such an abstract representation of q is useful for accommodating applications where q may include the non-numerical features relevant to the the buyer's decisions (e.g., the brand and production time of a used car). Since q is a random variable, $\beta(q)$ also has a probability distribution. For ease of presentation, we make a mild technical assumption that the distribution of β does not have any point mass. However, our analysis applies similarly to the general case in which $\beta(q)$ contains point masses, but just with more complex notations (see online version for more details).

With slight abuse of notation, let $v(t)$ denote the buyer's expected utility for action 1 under his prior beliefs about q, namely, when no information is purchased. That is,

$$v(t) = \int_{q\in Q} v(q,t)g(q)\mathrm{d}q. \tag{1}$$

Mechanism Space and the Revelation Principle

To maximize revenue, the the seller can design arbitrary mechanisms with possibly multiple rounds of interactions with the buyer. The task of designing a revenue-maximizing mechanism can be intractable unless a well-defined and general mechanism space is specified. Prior work of (Bergemann, Bonatti, and Smolin 2018) restricts to the sale of experiments via only a single-round of information revelation. In this work, we consider a richer design space of mechanisms, in which the the seller is also allowed to contract the realized experiment outcomes (i.e., signals) and moreover, multiple rounds of information revelation and payments are allowed as well. Specifically, we consider the following set of *sequential mechanisms*.[6]

Definition 1 (**Sequential Mechanisms**). *A sequential mechanism is a mechanism that results in a finite extensive-form game between the seller and the buyer. Formally, let $C(n)$ be the set of all children nodes of node n. Then each non-leaf node n in the game tree is one of the following three types:*

- Transfer node*, which is associated with a (possibly negative) monetary transfer $p(n)$ to the seller and has a single child node.*
- Seller node *that reveals information. Any seller node associates each state of nature q with a distribution over $C(n)$, prescribing the probabilities of moving to its children nodes. That is, there is a function $\psi_n : Q \times C(n) \mapsto$*

[5] This implies that the type's density function $f(t) > 0, \forall t \in T$.

[6] This general class of mechanisms was first introduced and studied by (Babaioff, Kleinberg, and Paes Leme 2012), and was called the generic interactive protocols in their work.

$[0,1]$ *for each seller node n with $\sum_{c \in C(n)} \psi_n(c,q) = 1, \forall q \in Q$. Thus, a child node c carries information about q.*
- Buyer node*, which corresponds to an arbitrary set of buyer choices with every choice leading to a child node.*

The buyer's final decision of taking the active or passive action is made *after* the information selling process, and thus is not modelled in the above sequential mechanisms. Therefore, at each seller node, the seller's action is to choose a message to send to the buyer which determines the child node the game will move to; the buyer node has the similar functionality. Note that the mechanism is voluntary and the buyer is free to leave the mechanism at any stage.

When designing the revenue-optimal mechanism for selling physical goods, the celebrated revelation principle (Myerson 1979; Gibbard 1973) enables us to without loss of generality focus only on truthful and direct mechanisms. However, when selling information, sequential mechanisms can bring strictly more revenue than one-round mechanisms. We show that our setting admits a stronger revelation principle that allows us to consider w.l.o.g. the set of truthful, direct and one-round mechanisms.

To describe the space of *one-round mechanisms*, we need the notion of *experiments*, which formalize the way a the seller reveals information. Given a set of possible signals Σ, an experiment $\pi : Q \rightarrow \Delta_\Sigma$ is a mapping from the state q to a distribution over the signals in Σ. Such an experiment can be mathematically described by $\{\pi(\sigma|q)\}_{q \in Q, \sigma \in \Sigma}$ where $\pi(\sigma|q)$ is the probability of sending signal σ conditioned on state q. After observing signal σ, the buyer infers posterior probability about any state q via standard Bayes updates:

$$g(q|\sigma) = \frac{\pi(\sigma|q) \cdot g(q)}{\int_{q' \in Q} \pi(\sigma|q') \cdot g(q') \mathrm{d}q'} = \frac{\pi(\sigma|q) \cdot g(q)}{\mathbf{E}_{q' \sim G}[\pi(\sigma|q')]} \quad (2)$$

Consequently, conditioned on signal σ, if a buyer of type t takes the active action, his expected utility is $\int_{q \in Q} v(q,t) g(q|\sigma) \mathrm{d}q$.

Different experiments reveal different amount of information to the buyer, and thus are of different values. A *one-round mechanism* is a menu of experiments and prices that results in a single-round of interaction between the seller and the buyer.

Definition 2 (**One-round Mechanisms**). *A one-round mechanism $\mathcal{M}$, described by a menu $\{(p_t, \pi_t)\}_{t \in T}$, proceeds as follows:*

1. *The buyer is asked to report (possibly untruthfully) his type t;*
2. *The seller charges the buyer p_t;*
3. *The seller reveals information about q according to experiment π_t.*

A one-round mechanism can clearly be represented as a special sequential mechanism, with the 3 steps corresponding to a buyer node, followed by a transfer node, and then followed by a seller node. Though sequential mechanisms can generally contract experiment outcomes (when a seller node is followed by transfer nodes), any one-round mechanism only prices the experiment π_t at price p_t but does not contract the experiment outcomes.

Let $U(t';t)$ denote the expected utility of a buyer with type t reporting type t', defined as

$$U(t';t) = \sum_{\sigma \in \Sigma} \max \left\{ \int_{q \in Q} v(q,t) \pi_{t'}(\sigma|q) g(q) \, \mathrm{d}q \ , \ 0 \right\} - p_{t'}$$

A one-round mechanism is said to be *incentive compatible*, if it is the buyer's best interest to report his type truthfully, i.e., $U(t;t) \geq U(t';t), \forall t, t' \in T$. The following revelation principle shows that it is without loss of generality to consider direct, incentive compatible mechanisms and one-round in our model.

Lemma 1 (**Revelation Principle**). *For any sequential mechanism $\mathcal{M}$, there exists a direct, incentive compatible and one-round mechanism that achieves the same expected revenue as $\mathcal{M}$.*

Standard revelation principle argument implies that the seller can w.l.o.g incentivize truthful type report at the beginning. To prove Lemma 1, the non-trivial part is to argue that a single-round of payment and information revelation suffice. This is a consequence of our independence assumption between state q and buyer type t, such that it allows us to simply combine all steps of information revelation as a single experiment and combine all payments as a single upfront payment. A formal proof can be found in our arXiv version. Notably, the proof of Lemma 1 relies crucially on the independence of state q and buyer type t. Fundamentally, this is because with correlation among the buyer type and state, a buyer type t, if misreporting t', will perceive a different expected payment as the $p_{t'}$ perceived by the buyer type t' since t and t' hold different belief about q and thus the expected payments w.r.t. each signal realization (see the proof for more illustration). Next, we further simplify the mechanism design space. First, we show in Lemma 2 that it is without loss of generality to consider mechanisms with non-negative payments. While this result is intuitive, we point out that it does not trivially hold. In fact, when q and t are correlated, the full-surplus-extracting sequential mechanism of (Babaioff, Kleinberg, and Paes Leme 2012) may have to use *negative* payments. The proof of this lemma is showed in our arXiv version.

Lemma 2 (**Non-Negative Payments**). *There exists an optimal IC, IR and one-round mechanism in which $p_t \geq 0$ for all $t \in T$.*

Second, the following known result of (Bergemann, Bonatti, and Smolin 2018) shows that when pricing experiments, we can without loss of generality price *responsive* experiments, in which each signal leads to a unique buyer best response action. From this perspective, each signal in a responsive experiment can be viewed as an *obedient* action recommendation.

Lemma 3 ((Bergemann, Bonatti, and Smolin 2018)). *The outcome of any mechanism can be obtained by using responsive experiments.*

Formulating the Optimal Pricing Problem

Based on the above simplification of the design space, we now formulate the mechanism design problem. We start by introducing (functional) variables to describe a one-round mechanism with responsive experiments. We will think of the payment in the menu $\mathcal{M}$ as a function $p(t)$ of buyer types t. Since the buyer has two possible actions, any responsive experiment π_t for buyer type t only needs two signals. With slight abuse of notation, we use function $\pi(q,t) \in [0,1]$ to denote the probability of sending signal $\texttt{active}$ (interpreted as an obedient recommendation of the active action), conditioned on state realization q. Naturally, $[1-\pi(q,t)]$ is the probability of sending signal $\texttt{passive}$ conditioned on state q. Our goal is to derive a feasible menu — represented by functions $\pi^*(q,t)$ and $p^*(t)$ — that maximizes the seller's revenue.

$$\text{Seller Revenue:} \quad \max_{\pi,p} \int_{t\in T} f(t)p(t)\,\mathrm{d}t.$$

Note that this is a *functional optimization* problem since both $\pi(q,t), p(t)$ are functional variables that depend on continuous variable $t\in[t_1,t_2](=T)$ and abstract variable q from a measurable set Q. The remainder of this section is devoted to formulating constraints on $\pi(q,t), p(t)$ according to Lemma 1, 2 and 3.

Obedience constraints. Lemma 3 shows that any responsive experiment only needs to have two signals which make obedient recommendation of the active and passive action, respectively. This poses two constraints on the function $\pi(q,t)$:(1) $\int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q \geq 0, \forall t \in T$; (2) $\int_{q\in Q}[1-\pi(q,t)]v(q,t)g(q)\,\mathrm{d}q \leq 0, \forall t \in T$. The first constraint above ensures that when signal $\texttt{active}$ is sent to buyer type t, the buyer's expected value $\frac{\mathbf{E}_{q\sim G}[\pi(q,t)v(q,t)]}{\mathbf{E}_{q\sim G}[\pi(q,t)]}$ for taking the active action is indeed at least 0, which is the expected value of taking the passive action. Similarly, the second constraint ensures the obedience of the $\texttt{passive}$ signal. Slightly manipulating the second constraint above, we obtain $\int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q \geq \int_{q\in Q}v(q,t)g(q)\,\mathrm{d}q = v(t)$, where $v(t)$ defined in Equation (1) is the buyer's a priori expected value of the active action. Therefore, we can conveniently summarize the obedience constraint as follows:

$$\text{Obedience:} \quad \int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q \geq \max\{0, v(t)\}, \forall t\in T \tag{3}$$

Individual rationality (IR) constraints. Since the the buyer gets utility 0 from the passive action, the expected utility of buyer type t, if he reports his type *truthfully* and follows the seller's obedient recommendation, is

$$\begin{aligned} u(t) &= \mathbf{E}_{q\sim G}[\pi(q,t)v(q,t)] - p(t) \\ &= \int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q - p(t) \end{aligned} \tag{4}$$

where the first term is the value from his decision making assisted by the seller's information and the second term is the payment to the seller. To ensure the buyer's participation in the mechanism, the following individual rationality (IR) constraint is required:

$$\text{IR:} \quad \int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q - p(t) \geq \max\{0, v(t)\}, \forall t\in T \tag{5}$$

where the right-hand side is the buyer's expected utility of not participating in the mechanism and simply takes the best action according to his prior belief about q. Interestingly, since the payment function is always non-negative according to Lemma 2, the IR constraint (5) turns out to imply the obedience constraint (3).

The buyer surplus $s(t)$ — the additional utility gain of participating in the mechanism — as a function of the buyer type t is defined as follows:

$$s(t) = \int_{q\in Q}\pi(q,t)v(q,t)g(q)\,\mathrm{d}q - p(t) - \max\{v(t), 0\} \tag{6}$$

The IR Constraint (5) is equivalent to non-negative surplus.

Incentive compatibility (IC) constraints. The derivation of the IC constraints turns out to be more involved. IC requires that when reporting truthfully, a buyer of type t should obtain a higher utility than misreporting any other type t'. This turns out to require some analyses since when a buyer of type t misreports type t', the resulting experiment $\{\pi(q,t')\}_{q\in Q}$ may not be obedient for t any more, leading to non-linearity in the IC constraints. Specifically, upon receiving signal $\texttt{active}$, the expected value of the active action for a type-t buyer misreporting t' is

$$\begin{aligned} V_a(t';t) &= \int_{q\in Q}\pi(q,t')v(q,t)g(q)\,\mathrm{d}q \\ &= \int_{q\in Q}\pi(q,t')\alpha(q)[t+\beta(q)]g(q)\,\mathrm{d}q \end{aligned}$$

Since $\pi(q,t')$ may not be obedient for buyer type t, he will choose between active action and the passive action, leading to true expected value $\max\{V_a(t';t),0\}$ in this situation.

Similarly, upon receiving signal $\texttt{passive}$, the buyer's value is the maximum between 0 and the following:

$$\int_{q\in Q}[1-\pi(q,t')]v(q,t)g(q)\,\mathrm{d}q = v(t) - V_a(t';t) \tag{7}$$

Combining both situations, the expected utility obtained by a buyer of type t from misreporting type t' is $\max\{V_a(t';t),0\} + \max\{v(t)-V_a(t';t),0\} - p(t')$. So the incentive compatibility constraint becomes the following:

$$u(t) \geq \max\{V_a(t';t),0\} + \max\{v(t)-V_a(t';t),0\} - p(t') \tag{8}$$

Such non-linear constraints are difficult to handle in general. Interestingly, it turns out that we can leverage previous results to reduce Constraint (8) to linear constraints on π, with some careful case analysis:

1. When $t > t'$, we have $V_a(t';t) \geq V_a(t';t') \geq 0$, where the first inequality is due to the assumption $\alpha(q) \geq 0$ and the second comes from the obedience constraint (3) for t'. In this case, the right-hand side of Constraint (8) becomes $V_a(t';t) + \max\{v(t) - V_a(t';t), 0\} - p(t')$, or equivalently $\max\{v(t), V_a(t';t)\} - p(t')$. Note that $u(t) \geq v(t) - p(t')$ is already implied by the IR constraint $u(t) \geq v(t)$ and the condition $p(t') \geq 0$. Therefore, the only non-redundant constraint in this case is $u(t) \geq V_a(t';t) - p(t')$.
2. When $t < t'$, we have $v(t) - V_a(t';t) \leq v(t') - V_a(t';t') \leq 0$ for similar reasons. In this case, the right-hand side of the above constraint becomes $\max\{V_a(t';t), 0\} - p(t')$. Again, $u(t) \geq -p(t')$ is already implied by the IR constraint $u(t) \geq 0$ and the condition $p(t') \geq 0$. Therefore, the only non-redundant constraint in this case is also $u(t) \geq V_a(t';t) - p(t')$.

To summarize, given the IR and non-negative payment constraints, the IC constraint can finally be reduced to the following:

$$\text{IC:} \quad \int_{q\in Q} \pi(q,t)v(q,t)g(q)\,\mathrm{d}q - p(t)$$
$$\geq \int_{q\in Q} \pi(q,t')v(q,t)g(q)\,\mathrm{d}q - p(t'), \forall t,t' \in T \tag{9}$$

Combined optimization problem. The derivation and simplification above ultimately lead to the following optimization problem, with functional variables $\pi(q,t), p(t)$:

$$\begin{array}{ll} \text{maximize} & \int_{t\in T} f(t)p(t)\,\mathrm{d}t \\ \text{subject to} & \int_{q\in Q} \pi(q,t)v(q,t)g(q)\,\mathrm{d}q - p(t) \\ & \geq \max\{0, v(t)\} \quad \forall t \in T \\ & \int_{q\in Q} [\pi(q,t) - \pi(q,t')]v(q,t)g(q)\,\mathrm{d}q \\ & -p(t) + p(t') \geq 0 \quad \forall t,t' \in T \\ & p(t) \geq 0, \quad \pi(q,t) \in [0,1] \end{array} \tag{10}$$

The Optimal Mechanism

In this section, we present the characterization of the optimal pricing mechanism. Mathematically, we derive an optimal solution in closed-form to the functional optimization problem (10). The optimal mechanism we obtain turns out to belong to the following category of *threshold mechanisms*.

Definition 3 (**Threshold Mechanisms**). *A mechanism (π, p) is called a threshold mechanism if it only uses threshold experiments. That is, there exists a function $\theta(t)$, such that for any $t \in [t_1, t_2]$,*

$$\pi(q,t) = \begin{cases} 1 & \text{if } \beta(q) \geq \theta(t) \\ 0 & \text{otherwise} \end{cases}.$$

In this case, $\pi(q,t)$ is fully described by the threshold function $\theta(t)$.

Note that the term "threshold" is only a property about the experiments and does not pose any constraint on the payment function $p(t)$. To formally present our mechanism, we will need the following notions of *lower*, *upper* and *mixed* virtual value functions.

Definition 4 (**Lower/Upper/Mixed Virtual Value function**). *For any type t with PDF $f(t)$ and CDF $F(t)$, the function $\underline{\phi}(t) = t - \frac{1-F(t)}{f(t)}$ is called the* lower virtual value function *and $\bar{\phi}(t) = t + \frac{F(t)}{f(t)}$ is called the* upper virtual value function. *Moreover, for any $c \in [0,1]$, $\phi_c(t) = c\underline{\phi}(t) + (1-c)\bar{\phi}(t)$ is called a* mixed virtual value function.

Any virtual value function is regular *if it is monotone non-decreasing in t.*

The lower virtual value function $\underline{\phi}(t)$ is precisely the virtual value function commonly used in classic mechanism design settings (Myerson 1981). We remark that while the upper and mixed virtual value function were not formally defined before, they have implicitly shown up in previous works and typically give rise when the IR constraints are binding at the largest type (e.g., (Eső and Szentes 2007b)). However, the specific formulation for the information selling problem allows us to characterize the optimal mechanism for much more general buyer utility functions (see more detailed comparison in the related work).

Ironing. When a virtual value function is irregular, we will need to apply the so-called "ironing" trick to make it monotone non-decreasing in t. (Myerson 1981) developed a procedure for ironing the lower virtual value function $\underline{\phi}(t)$. This procedure can be easily generalized to iron any function about the buyer type t, specifically, also to the three types of the virtual value functions defined above. For any virtual value function $\phi(t)$ (upper, lower or mixed), let $\phi^+(t)$ denote the ironed version of $\phi(t)$ obtained via the standard ironing procedure of (Myerson 1981) (for completeness, we give a formal description of this ironing procedure in the arXiv version).[7]

If a virtual value function $\phi(t)$ is already non-decreasing, it remains the same after ironing, i.e., $\phi^+(t) = \phi(t), \forall t$. With $\phi_c(t) = c\underline{\phi}(t) + (1-c)\bar{\phi}(t)$, the following useful properties of the ironed mixed virtual value functions will be needed for proving our main result (and may also be of independent interest in general). Their proofs are technical and are deferred to the arXiv version.

Lemma 4 (**Useful Properties of Ironed Mixed Virtual Values**).

1. *For any $0 \leq c < c' \leq 1$, $\phi^+_c(t) \geq \phi^+_{c'}(t)$ for any t;*
2. *For any $c \in [0,1]$, let t_c be the buyer type such that $F(t_c) = c$. Then we have $\phi^+_c(t) \leq t, \forall t \leq t_c$ and $\phi^+_c(t) \geq t, \forall t \geq t_c$. This also implies $\underline{\phi}^+(t) < t < \bar{\phi}^+(t), \forall t \in (t_1, t_2)$.*

Notably, the second property above also implies that $\phi^+_c(t_c) = t_c$ always holds.

We will be readily prepared to state the optimal mecha-

[7] For techniques to iron a general function, we refer the reader to a recent work by (Toikka 2011).

nism after introducing the following two quantities:

$$V_L = \max\{v(t_1), 0\} + \int_{t_1}^{t_2} \int_{q:\beta(q) \geq -\underline{\phi}^+(x)} g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}x, \quad (11)$$

$$V_H = \max\{v(t_1), 0\} + \int_{t_1}^{t_2} \int_{q:\beta(q) \geq -\bar{\phi}^+(x)} g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}x, \quad (12)$$

where $\bar{\phi}^+(x)$ and $\underline{\phi}^+(x)$ are the ironed upper and lower virtual value functions, respectively. Note that Lemma 4 implies $-\underline{\phi}^+(x) \geq -\bar{\phi}^+(x)$ and consequently $V_L \leq V_H$ since $g(q)\alpha(q)$ is always non-negative and thus V_L integrates over a smaller region.

Our main result is then summarized in the following theorem.

Theorem 1 (Characterization of an Optimal Mechanism).

1. *If $v(t_2) \leq V_L$, the threshold mechanism with threshold function $\theta^*(t) = -\underline{\phi}^+(t)$ and the following payment function represents an optimal mechanism:*

$$p^*(t) = \int_{q \in Q} \pi^*(q,t) g(q) v(q,t)\, \mathrm{d}q - \int_{t_1}^{t} \int_{q \in Q} \pi^*(q,x) g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}x.$$

where π^* is determined by $\theta^*(t)$ as in Definition 3. Moreover, $p^*(t)$ is monotone non-decreasing for $t \in [t_1, t_2]$.

2. *If $v(t_2) \geq V_H$, the threshold mechanism with threshold function $\theta^*(t) = -\bar{\phi}^+(t)$ and the following payment function represents an optimal mechanism:*

$$p^*(t) = \int_{q \in Q} \pi^*(q,t) g(q) v(q,t)\, \mathrm{d}q + \int_{t}^{t_2} \int_{q \in Q} \pi^*(q,x) g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}x - v(t_2),$$

where π^* is determined by $\theta^*(t)$ as in Definition 3. Moreover, $p^*(t)$ is monotone non-increasing for $t \in [t_1, t_2]$.

3. *If $V_L < v(t_2) < V_H$, let $c \in (0,1)$ be a constant that satisfies*

$$\int_{t_1}^{t_2} \int_{q:\beta(q) \geq -\phi_c^+(t)} g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}t = v(t_2),$$

where $\phi_c^+(t)$ is the ironed version of the mixed virtual value function $\phi_c(t)$. Then the threshold mechanism with threshold function $\theta^*(t) = -\phi_c^+(t)$ and the following payment function represents an optimal mechanism:

$$p^*(t) = \int_{q \in Q} \pi^*(q,t) g(q) v(q,t)\, \mathrm{d}q - \int_{t_1}^{t} \int_{q \in Q} \pi^*(q,x) g(q)\alpha(q)\, \mathrm{d}q \mathrm{d}x.$$

Moreover, $p^*(t)$ is monotone non-decreasing in t when $F(t) \leq c$ and monotone non-increasing when $F(t) > c$.

Let $\bar{t}$ satisfy $v(\bar{t}) = 0$. In all cases above, the buyer surplus function $s(t)$ is convex and monotone non-decreasing when $t \leq \bar{t}$, but immediately transitions to be convex and monotone non-increasing when $t \geq \bar{t}$.

The Power of Information Discrimination

The above example shows that the optimal mechanism features information discrimination, i.e., reveals different information to different buyer types, which then leads to price discrimination. One might wonder how well a mechanism can perform if information discrimination is not allowed. Our following proposition shows that in this case, the optimal mechanism is to simply post a uniform price and then reveal full information to any buyer who is willing to pay.

To describe the mechanism, we introduce a notation $e(t)$ that captures the value of full information for any buyer with type t:

$$e(t) = \int_{q \in Q} \max\{v(q,t)\,,\, 0\} g(q)\, \mathrm{d}q - \max\left\{ \int_{q \in Q} v(q,t) g(q)\, \mathrm{d}q\,,\, 0 \right\} \quad (13)$$

That is, $e(t)$ equals the additional value buyer type t obtains by fully observing q. We have the following lemmas and their proofs can be found in the arXiv version.

Lemma 5. *If information discrimination is not allowed, then the optimal mechanism is to charge the Myerson's reserve price r^* with respect to value $e(t)$, i.e., $r^* = \mathrm{argmax}_r\, [r \cdot \mathbf{Pr}_{t \sim F}(e(t) \geq r)]$, and then reveal full information to any buyer who pays.*

Lemma 6. *If distribution of $e(t)$ has monotone hazard rate (with randomness inherited from $t \sim F$), then we always have $\frac{RevSingle^*}{Rev^*} \geq \frac{1}{e}$.*

Proof of the Main Theorem

In this section, we prove Theorem 1. Due to space limit, we will only provide a complete proof for Case 3. The core idea for proving Case 1 and 2 is similar. We thus defer them to our arXiv version, respectively. The proof has two major steps: (1) characterizing useful properties of (any) feasible mechanisms; (2) leveraging the properties to derive the optimal mechanism. While the first step is also based on the analysis of the IC constraints as in classic mechanism design, the conclusions we obtain are quite different since our problem's constraints are different. Significantly deviating from the Myersonian approaches for classic mechanism design is our second main step, which arguably is much more involved due to additional constraints that we have to handle (this is also reflected in the more complex format of our optimal mechanism).

Useful Properties of Feasible Mechanisms

Define *feasible* mechanisms as the set of mechanisms (π, p) that satisfy all the constraints of program (10) (but not necessarily maximizing its objective). We first characterize the space of feasible mechanisms. To describe our characterization, it is useful to introduce the following quantity.

$$P_\pi(t) = \int_{q\in Q} \pi(q,t)g(q)\alpha(q)\,\mathrm{d}q \tag{14}$$

Note that $P_\pi(t)$ can be interpreted as the expected *weighted* probability (with weight $\alpha(q)$) of being recommended the active action 1. The following lemma summarizes our characterization. Its proof is showed in our arXiv version.

Lemma 7 (Characterization of Feasible Mechanisms). *A mechanism (π, p) with non-negative payments is feasible if and only if it satisfies the following constraints:*

$$P_\pi(t) \text{ is monotone non-decreasing in } t \tag{15}$$

$$u(t) = u(t_1) + \int_{t_1}^{t} P_\pi(x)\,\mathrm{d}x, \forall t\in T \tag{16}$$

$$u(t_2) \geq v(t_2), \;\; u(t_1) \geq 0 \tag{17}$$

$$p(t) \geq 0, \;\forall t \in T \tag{18}$$

Note that condition (15) is analogous to Myerson's allocation monotonicity condition in the auction design problem, but also differs in the sense that the value of an item in auction design only depends on the buyer type t with no weight associated to it. In information selling, the value of taking the active action will depend on the utility coefficient $\alpha(q)$.

Next we characterize the buyer's surplus $s(t) = u(t) - \max\{0, v(t)\}$, as expressed in Equation (6), from participating in the information selling mechanism. Recall that, with only the prior information, a buyer of type t has expected utility $v(t) = \int_{q\in Q} v(q,t)g(q)\,\mathrm{d}q$ for the active action. Since $v(q,t)$ is monotone non-decreasing in t, we know that $v(t)$ is also monotone non-decreasing. Let $\bar{t}$ be any buyer type at which $v(t) = 0$. The following lemma characterize how the buyer's surplus changes as a function of his type.

Lemma 8. *Let $\bar{t}$ be any buyer type such that $v(\bar{t}) = \int_{q\in Q} v(q,\bar{t})g(q)\,\mathrm{d}q = 0$. In any feasible mechanism (π, p) with non-negative payments, the buyer's surplus $s(t)$ is monotone non-decreasing for $t \in [t_1, \bar{t}]$ and monotone non-increasing for $t \in [\bar{t}, t_2]$.*[8]

Proof. When $t \leq \bar{t}$, we have $v(t) \leq 0$. Therefore, without participating in the mechanism to purchase additional information, the buyer will get maximum utility 0 by taking the passive action. So his surplus for participation is

$$s(t) = u(t) = u(t_1) + \int_{t_1}^{t} P_\pi(x)\,\mathrm{d}x$$

by the utility identify in Equation (16). Since $u(t_1) \geq 0$ and $P_\pi(x) \geq 0$, it is easy to see that $s(t)$ is non-negative and monotone non-decreasing in t.

[8] $\bar{t}$ can be any one of them if there are multiple t such that $v(t) = 0$. If no $\bar{t} \in [t_1, t_2]$ makes $v(\bar{t}) = 0$, then either $v(t) < 0$ or $v(t) > 0$ for any $t \in T$ and in this case $s(t)$ is monotone within T.

When $t \geq \bar{t}$, we have $v(t) \geq 0$. So the buyer's maximum utility is $v(t)$ without participating in the information selling mechanism. We thus have

$$\begin{aligned}
s(t) =& u(t) - v(t) \\
=& \left[u(t_1) + \int_{t_1}^{t}\int_{q\in Q} \pi(q,x)\alpha(q)g(q)\,\mathrm{d}q\mathrm{d}x\right] \\
&- \left[\int_{q\in Q} \alpha(q)[t+\beta(q)]g(q)\,\mathrm{d}q\right] \\
=& \left[u(t_1) + \int_{t_1}^{t}\int_{q\in Q} \pi(q,x)\alpha(q)g(q)\,\mathrm{d}q\mathrm{d}x\right] \\
&- \left[\int_{t_1}^{t}\int_{q\in Q} \alpha(q)g(q)\,\mathrm{d}q\mathrm{d}x + v(t_1)\right] \\
=& u(t_1) - v(t_1) + \left[\int_{t_1}^{t}\int_{q\in Q} [\pi(q,x)-1]\alpha(q)g(q)\,\mathrm{d}q\mathrm{d}x\right].
\end{aligned}$$

Since $\pi(q,x) - 1 \leq 0$ and $\alpha(q)g(q) \geq 0$, we thus have that $s(t)$ is monotone non-increasing in t. Notably, $s(t) \geq s(t_2) = u(t_2) - v(t_2) \geq 0$ by inequality (17). ☐

After holding those properties, we are ready to proof the main theorem case by case with side lemmas. The proof is intricate and technical as we need to carefully balance the surplus curve. Thus, we refer readers to our arXiv version section 4.2 to gain a further picture.

References

Babaioff, M.; Kleinberg, R.; and Paes Leme, R. 2012. Optimal Mechanisms for Selling Information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, 92–109. New York, NY, USA: ACM. ISBN 978-1-4503-1415-2.

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

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

Chen, Y.; Xu, H.; and Zheng, S. 2020. Selling Information Through Consulting. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2412–2431. SIAM.

Eső, P.; and Szentes, B. 2007b. The price of advice. *RAND Journal of Economics*, 38(4): 863–880.

Gibbard, A. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4): 587–601.

Hörner, J.; and Skrzypacz, A. 2016. Selling information. *Journal of Political Economy*, 124(6): 1515–1562.

Myerson, R. 1981. Optimal Auction Design. *Mathematics of Operations Research*, 6(1): 58–73.

Myerson, R. B. 1979. Incentive Compatibility and the Bargaining Problem. *Econometrica*, 47(1): 61–73.

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