Discussion Paper No. 1294

ISSN (Print) 0473-453X
ISSN (Online) 2435-0982

# DYNAMIC PRICING OF INFORMATION: BELIEF DIVERGENCE AND SURPLUS EXTRACTION

Chia-Hui Chen
Junichiro Ishida
Wing Suen

October 2025

The Institute of Social and Economic Research
The University of Osaka
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

# Dynamic Pricing of Information: Belief Divergence and Surplus Extraction*


CHIA-HUI CHEN
Kyoto University

JUNICHIRO ISHIDA
University of Osaka

WING SUEN
University of Hong Kong


October 20, 2025


*Abstract.* We consider a screening problem in which a data seller offers a dynamic payment schedule for a sequence of experiments to a privately informed buyer. Different buyer types face different expected costs for the same payment schedule. This payment gap can be optimized to reduce information rent. Dynamic mechanisms strictly increase seller revenue compared to the optimal static mechanism, and may even extract full surplus under some conditions. We obtain a full characterization of optimal dynamic mechanisms, which can take the simple form of a binary experiment at each stage. Payments are back-loaded and experiments become progressively more informative over stages.



*Keywords.* monopoly screening; dynamic experiment sequence; interim individual rationality constraint; belief divergence; full surplus extraction; payment gap

*JEL Classification.* D80; D82



*We thank Jin Li, Bart Lipman, Xianwen Shi, and Takuro Yamashita for helpful discussions on this topic.

## 1. Introduction

A recurrent theme in the theory of monopoly screening under private information (Mussa and Rosen, 1978; Maskin and Riley, 1984) is that the seller may distort the quantity or quality of the product offered to low-demand buyers in order to make it less attractive for high-demand buyers to mimic low-demand types, thus reducing their information rent. A recent contribution by Bergemann et al. (2018) shows that this logic extends to a data seller who designs and sells experiments to potential buyers with heterogeneous private beliefs. They characterize optimal static mechanisms and find that the seller may provide buyers with low demand for information an imperfect experiment, even though it is costless to produce fully informative ones. We extend the framework of Bergemann et al. (2018) by allowing the seller to offer a dynamic mechanism in which she sells a sequence of experiments over multiple stages. Because the underlying state is not changing, any sequence of experiments can be collapsed into one single static experiment that is informationally equivalent to the dynamic sequence. Therefore, in an environment where buyers have quasi-linear preferences, it is not immediately obvious why and how a dynamic mechanism can improve on static ones. Nevertheless we show that the seller always strictly gains from selling experiments over stages as long as there remains surplus to be extracted from trade. The gains from offering a dynamic mechanism can be substantial, even enabling the seller to extract full surplus under some conditions.

To fix ideas, consider a simple example in which a buyer wants to choose an action $a \in \{a_0, a_1\}$ to match the state $\omega \in \{\omega_0, \omega_1\}$. His payoff is \$100 if the action matches the state ($a_i = \omega_i$), or zero otherwise. Before entering the market, the buyer privately observes a signal and updates his belief. Specifically, a "low" type has belief 0.75 (that the state is $\omega_1$) and a "high" type has belief 0.3. The seller is uninformed about the buyer type and believes that the two types are equally likely. Note that the high type is more uncertain about the state; the willingness-to-pay for the fully informative experiment is therefore higher for the high type (\$30) than for the low type (\$25).[1] The maximum surplus the seller can extract in this environment is $0.5(\$30 + \$25) = \$27.5$.

Suppose the seller offers a binary experiment yielding two signals, say "up" and "down," which cause the buyer to update his belief upward and downward, respectively. The seller

[1]We say that an experiment is fully informative if it reveals the true state with probability 1. Since the fully informative experiment always allows the buyer to choose the right action, its value to a buyer with belief $\theta$ is $\$100(1 - \max\{\theta, 1 - \theta\})$.

offers a menu of experiments to maximize her revenue. In this particular setting, the optimal static mechanism is to offer Menu A:

- For the high type, the fully informative experiment at price \$30;
- For the low type, experiment $e^A$ at price $\$200/9 \approx 22.22$,

where the conditional distribution of signals given each state under $e^A$ is specified as below.

| | Experiment $e^A$ | |
|---|---|---|
| | up signal | down signal |
| state $\omega_1$ | 1 | 0 |
| state $\omega_0$ | 1/9 | 8/9 |

The price \$200/9 for experiment $e^A$ is chosen to bind the low type's individual rationality constraint. The seller optimally distorts $e^A$ to make it unattractive to the high type. Indeed, the high type's information rent from mimicking the low type is

$$100\left[0.3 + 0.7 \cdot \frac{8}{9} - (1 - 0.3)\right] - \frac{200}{9} = 0.$$

The seller's expected revenue from this mechanism is $\$235/9 \approx 26.11$.

Now suppose the seller can devise and commit to a two-stage mechanism. Consider Menu B:

- For the high type, the fully informative experiment at price \$30;
- For the low type, experiment $e_1^B$ at price $p_1^B = 1835/117 \approx \$15.68$, followed by
  - experiment $e_2^B$ at price $p_2^B = 900/65 \approx \$13.85$ if $e_1^B$ yields a down signal,
  - no further experiment at no cost if $e_1^B$ yields an up signal,

where $e_1^B$ and $e_2^B$ are specified as below.

| | Experiment $e_1^B$ | | Experiment $e_2^B$ | |
|---|---|---|---|---|
| | up signal | down signal | up signal | down signal |
| state $\omega_1$ | 2/3 | 1/3 | 1 | 0 |
| state $\omega_0$ | 1/9 | 8/9 | 0 | 1 |

The sequence of experiments offered to the low type in Menu B would generate the same distribution of final actions in each state (i.e., action $a_1$ with probability 1 in state $\omega_1$, and action $a_0$ with probability 8/9 in state $\omega_0$) as that generated by the static experiment

$e^A$ in Menu A. Thus, this experiment sequence is informationally equivalent to $e^A$, and the willingness-to-pay for this experiment sequence remains at \$200/9 for both types. The prices $p_1^B$ and $p_2^B$ bind the low type's individual rationality constraint at the initial stage, as they satisfy

$$p_1^B + \left[0.75 \cdot \frac{1}{3} + 0.25 \cdot \frac{8}{9}\right] p_2^B = \frac{200}{9},$$

where the term in brackets reflects the low type's belief that the first-stage experiment $e_1^B$ will yield a down signal.[2] By construction, Menu B will produce the same expected revenue to the seller as Menu A. The key difference, however, is that the high type's incentive compatibility constraint becomes slack in Menu B. His expected payment for the experiment sequence offered to the low type is

$$p_1^B + \left[0.3 \cdot \frac{1}{3} + 0.7 \cdot \frac{8}{9}\right] p_2^B = \frac{3005}{117} > \frac{200}{9}.$$

This payment gap—the difference between the two types' expected payments for the same experiment sequence—comes from the fact that the high type (who has a stronger belief about state $\omega_0$) attaches a larger probability that the first-stage experiment $e_1^B$ will yield a down signal.

As a two-stage mechanism, it is clear that Menu B is not optimal because a relevant incentive constraint is slack. There are two ways to strictly improve upon Menu B. First, the experiment sequence in Menu B is equally informative as $e^A$. By offering an experiment sequence more informative than $e^A$ and charging a higher price for it, the seller can obtain more revenue from the low type without violating the high type's incentive compatibility constraint, which holds with slack. Second, while Menu B leads to a positive payment gap, the optimal dynamic mechanism should back-load payments as much as feasible to *maximize* the payment gap.[3] We will provide an explicit characterization and construction of optimal dynamic mechanisms in this paper. In the current example, the optimal two-stage mechanism turns out to be Menu C:[4]

[2] If the high type were to mimic the low type and obtain a down signal from $e_1^B$ at the first stage, his belief would be revised downward to $9/65 \approx 0.1385$. The second-stage price $p_2^B \approx \$13.85$ prevents the high type from double deviation, in which he pays for $e_1^B$ but refuses to pay for $e_2^B$.

[3] For example, raising the probability of the down signal in state $\omega_1$ under the first-stage experiment to slightly above $1/3$ would raise the high type's belief upon receiving a down signal to slightly higher than $9/65$, allowing the seller to charge a higher price $p_2^B$ for the second-stage experiment. The tradeoff is that this change would reduce the difference between the two types' assessments of the probability of receiving a down signal. The optimal mechanism has to balance this tradeoff.

[4] For Menu C, we use decimals because the solution is only approximate. The exact solution involves

- For the high type, the fully informative experiment at price \$30;
- For the low type, experiment $e_1^C$ at price $p_1^C = \$14.96$, followed by
  - experiment $e_2^C$ at price $p_2^C = \$16.34$ if $e_1^C$ yields a down signal,
  - no further experiment at no cost if $e_1^C$ yields an up signal,

where $e_1^C$ and $e_2^C$ are specified as below.

| | Experiment $e_1^C$ | | Experiment $e_2^C$ | |
|---|---|---|---|---|
| | up signal | down signal | up signal | down signal |
| state $\omega_1$ | 0.5556 | 0.4444 | 1 | 0 |
| state $\omega_0$ | 0.0244 | 0.9756 | 0 | 1 |

Menu C gives rise to an expected revenue of \$27.20 to the seller. The experiment sequence in Menu C is more informative than $e^A$ in Menu A. Payments are also more back-loaded, and the payment gap is maximized. These two factors account for the strict improvement of Menu C over the optimal static mechanism.

This proposed mechanism works as it effectively enables the seller to charge different prices to different types for the same item. This method of price discrimination is usually considered infeasible in the standard screening setting where the item for sale is a physical object. In contrast, in the context of trading information, differences in initial beliefs held by different types create both a gap between their willingness-to-pay for the product (an experiment sequence), and a gap between their expectations about whether they have to pay for further experiments or not. The seller can exploit this feature by designing the experiment sequence in such a way that the high type assigns a higher probability to reaching the second stage. Back-loading the payments makes the low type's experiment sequence less attractive to the high type, thus providing an additional channel (beyond distorting the nature of the product) of relaxing the incentive constraints arising from private information.

The logic of back-loading payments to enlarge the expected payment gap suggests that the seller may do better than offering two-stage mechanisms by adding even more stages. A complete analysis of dynamic mechanisms with multiple stages is, however, challenging because the space of possible mechanisms is very large and could expand exponentially with the number of stages. To complicate matters, in the dynamic setting with binary actions, it is not immediately clear whether restricting attention to binary experiments is

irrational numbers.

without loss of generality. In a static mechanism, if an experiment designed for some type yields multiple signals that lead to the same action, it is always weakly profitable to merge them into a single one, as it does not affect the informativeness of the experiment for that type, while it can only reduce the informativeness for other types. This argument needs to be refined here because the expected payment for each type in a dynamic mechanism depends on the whole sequence of beliefs the buyer expects to follow, which is inevitably affected when some signals are merged together along the sequence.

Despite these difficulties, we establish that it is without loss of generality to focus on binary experiments. In an optimal mechanism, the experiment sequence terminates after an up signal, while it continues after a down signal until the belief reaches 0. These properties allow us to substantially reduce the class of mechanisms to be considered. We show that optimal dynamic mechanisms have the following features:

- Payments are back-loaded, with a constant price after the second stage, while experiments become progressively more informative over stages; distortions, if any, are front-loaded and appear at the first stage.
- Full surplus extraction is feasible under some conditions; and when it is feasible, it can be implemented by a finite-stage mechanism.
- Adding another stage to a $T$-stage mechanism strictly improves the seller's revenue unless the mechanism already achieves full surplus extraction.

Given that full surplus extraction can never be achieved by a static mechanism, the final point implies that an optimal $T$-stage mechanism is strictly better than the optimal static mechanism for any $T \geq 2$, underscoring the value of dynamic mechanisms in this trading environment.

*Literature*. This paper builds on the model of information pricing by a data seller using a mechanism-design approach, pioneered by Bergemann et al. (2018); see also Bergemann and Bonatti (2019) for a survey of the literature on the market for information. In an online appendix to their paper, Bergemann et al. (2018) suggest the possibility of extending their approach to study dynamic mechanisms, and provide an example showing that a two-stage mechanism can improve on the optimal static mechanism. We take on this research direction to provide a full characterization of optimal dynamic mechanisms with binary states and actions, and show how the insights extend to more general environments. Section 4.2 provides a further discussion of how our results differ from those provided by Bergemann et al. (2018). From a broader perspective, we aim to contribute to the extensive literature on mechanism design by shedding new light on how sellers may extract surplus along a dynamic dimension. We highlight the back-loading of payments as the key instrument for relaxing incentive compatibility constraints, introducing an additional dimension to the designer's toolbox.

The way information disclosure occurs gradually over stages in our model is reminiscent of Hörner and Skrzypacz (2016), although the two model settings are very different. Most notably, the seller's decision in our model concerns not only how much information to disclose at each stage (its accuracy), but also along which state to disclose more information (its direction). Buyer types are horizontally differentiated in our setting, with no obvious direction that the seller seeks to induce. In Hörner and Skrzypacz (2016), the informed party is the seller who is either competent or incompetent; types are thus vertically differentiated, where both types want to be perceived as competent. We provide further elaboration on this point in Section 4.3.

Several works analyze mechanism-design problems with information disclosure (Eső and Szentes, 2007; Li and Shi, 2017; Krähmer, 2020; Bergemann et al., 2022; Wei and Green, 2024). They consider various trading mechanisms in which the seller can reveal some additional information that affects the buyers' valuations for the seller's product. In our setting, the object for sale is information itself. A connection can also be made to the literature on information design with heterogeneous or privately informed receivers (Alonso and Camara, 2016; Kolotilin et al., 2017; Guo and Shmaya, 2019; Chan et al., 2019). Here, although the data seller designs experiments to control the flow of information to the buyer, the seller's goal is not persuasion; rather, it is to extract as much surplus as possible from the value she provides to different types of buyers. The possibility of full surplus extraction in trading environments under asymmetric information has been discussed extensively since Crémer and McLean (1985, 1988) and McAfee and Reny (1992).[5] Models in this strand of literature typically exploit the correlation of types across agents to induce truth-telling. Our reasoning is fundamentally different, as we consider a setting with a single buyer.

In addition, there is an extensive literature on trading among agents with heterogeneous beliefs (Sebenius and Geanakoplos, 1983; Morris, 1994; Coval and Thakor, 2005). In this strand of literature, two agents with different priors can potentially gain from betting on their beliefs. The logic of our construction is quite different, as buyers of different

[5]Krähmer (2020) shows that full surplus extraction is feasible in a mechanism-design environment with information disclosure. Dynamic versions of such models are analyzed by Liu (2018) and Noda (2019).

types never trade with one another in our setting. The data seller is at an informational disadvantage relative to the buyer, but leverages a dynamic mechanism to exploit belief differences across buyer types.

## 2. Model

*Environment.* Consider a model with a data seller and a buyer trading over $T \geq 2$ stages. For most of the analysis, we take $T$ to be finite and exogenously given, but deducing the value of the optimal $T$ for the seller is quite straightforward once we characterize the optimal mechanism for every $T$. Section 4.4 studies the limit of the optimal mechanism as $T$ goes to infinity.

There are two states, $\omega \in \Omega = \{\omega_0, \omega_1\}$, and the state does not change over the trading stages. At each stage $t = 1, \ldots, T$, the buyer either purchases an experiment from the seller to obtain more information about the state or takes an action $a \in A := \{a_0, a_1\}$, where an experiment $e = (S, \pi)$ specifies a set of signals $S$ and a mapping $\pi : \Omega \to \Delta S$ from the state to a distribution over signals. The payoffs are realized, and trading ends when the buyer takes an action.

*Information.* The buyer and the seller are endowed with a common prior $\theta^0 \in (0,1)$ that $\omega = \omega_1$. Before trading begins, the buyer observes a private signal about the state and forms an initial belief, which is either $\theta_L^1$ with probability $q$ or $\theta_H^1$ with probability $1-q$; in what follows, we often abbreviate the superscript and simply write $\theta_i^1 = \theta_i$ to avoid clutter. This gives rise to two possible *types*, denoted by $\theta_i$, $i = H, L$. It is important that the initial beliefs of the two types are derived from privately observed signals rather than from heterogeneous priors. In our model, the buyer is more informed than the seller, and the seller's mechanism will leverage such information.

At each stage $t$, if the buyer chooses to purchase an experiment, he updates his belief upon observing a signal $s \in S$ generated by the experiment. Let $\mathscr{H}_t$ denote the set of all histories at stage $t$, with $\mathscr{H}_1 = \emptyset$, and $\mathscr{H} := \cup_{t=2}^{T} \mathscr{H}_t$. The history at stage $t \geq 2$ is the sequence of publicly observed realized signals up to (but not including) that stage, and is generically denoted by $h \in \mathscr{H}_t$. For $h \in \mathscr{H}_t$, we often use $hs \in \mathscr{H}_{t+1}$ to denote the history in which signal $s$ is observed following $h$. Let $\mathscr{H}_{T+1}$ represent the set of histories at the end of stage $T$, after the realization of the stage-$T$ signal.

*Mechanism.* The seller offers a menu of contracts $(M_L, M_H)$, or a mechanism. For $i = L, H$, each contract $M_i = (E_i, P_i)$ consists of an experiment sequence $E_i$ and a payment schedule $P_i$ when the reported type is $\theta_i$. An experiment sequence $E_i = \{e_i^h\}_{h \in \mathscr{H}_1 \cup \mathscr{H}}$ is a contingent plan of experiments that specifies what experiment to be conducted next at history $h$, whereas a payment schedule $P_i = \{p_i^h\}_{h \in \mathscr{H}_1 \cup \mathscr{H}}$ with $p_i^h \geq 0$ specifies the additional payment to the seller for the next experiment to be conducted. According to our notation, $p_i$ and $e_i$ (without superscripts) refer to the stage-1 price and stage-1 experiment offered to type $i$ following the initial history. Also, let $M_i^h = (E_i^h, P_i^h)$ represent the continuation mechanism at $h \in \mathscr{H}$.

*Commitment.* As is standard in the mechanism-design approach, the seller can fully commit to any experiment sequence and payment schedule at the outset. We do not require, however, the buyer to have the same commitment power. Specifically, we assume that the buyer may drop out at any point if it is not in his best interest to follow the experiment sequence. This implies that any optimal mechanism must satisfy *interim individual rationality*, in the sense that at every history $h \in \mathscr{H}$, if the reported type is $\theta_i$, then the buyer has the incentive to comply with paying $p_i^h$ for the next experiment $e_i^h$ specified in $E_i$. The assumption that buyers can drop out at any point is especially realistic for information markets, where they often interact with data sellers just online and anonymously.

*Payoffs.* If the buyer chooses $a_k$ in state $\omega_j$, his payoff is $\alpha > 0$ if $k = j = 1$, $\beta > 0$ if $k = j = 0$, or 0 if $k \neq j$. Define $\sigma := \beta/(\alpha + \beta)$ as the belief at which the buyer is indifferent between the two actions. Let

$$v(\theta) = \begin{cases} \theta\alpha & \text{if } \theta \geq \sigma \\ (1-\theta)\beta & \text{if } \theta < \sigma, \end{cases}$$

which is the buyer's expected payoff when he takes an action at belief $\theta$. The seller maximizes the expected revenue collected from the buyer. There is no discounting across trading stages.

# 3. Analysis

## 3.1. Preliminaries

Let $e^h = (S^h, \pi^h)$ be the experiment offered at history $h \in \mathscr{H}_1 \cup \mathscr{H}$, where $\pi^h(s|\omega)$ is the probability of observing $s \in S^h$ in state $\omega \in \Omega$. Once we fix an experiment sequence, we can compute the probability of reaching history $h \in \mathscr{H}$ in state $\omega_j$. Denote this probability by $\gamma_j^h$ for $j = 0, 1$. According to this definition, for any $s \in S^h$, $\gamma_j^{hs} = \gamma_j^h \pi^h(s|\omega_j)$. Also, for $i = L, H$, let $\theta_i^h$ be type $i$'s interim belief at $h$, given by

$$\theta_i^h = \frac{\theta_i \gamma_1^h}{\theta_i \gamma_1^h + (1-\theta_i)\gamma_0^h}.$$

The value of information of an experiment sequence $E$, or simply the value of an experiment sequence, can be measured by the net payoff gain from the sequence and denoted as $VI(\theta; E)$. At any history $h \in \mathscr{H}_1 \cup \mathscr{H}$, the value of the continuation experiment sequence $E^h$ can be written as

$$VI(\theta_i^h; E^h) = \mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^h, E^h\right] - v(\theta_i^h), \tag{1}$$

where the expectation is taken with respect to the random posterior belief $\tilde{\theta}_i$ induced by $E^h$ at the time the buyer takes an action conditional on the interim belief $\theta_i^h$. Let

$$\phi_i^h(s) := \theta_i^h \pi^h(s|\omega_1) + (1-\theta_i^h)\pi^h(s|\omega_0)$$

be the probability that type $i$ observes $s \in S^h$ at history $h$. Note that $\phi_H^h(s) > \phi_L^h(s)$ for any signal realization $s$ that will lead to downward belief revision at history $h$. The expected gross payoff, $\mathbb{E}[v(\tilde{\theta}_i)|\theta_i^h, E^h]$, is defined recursively by

$$\mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^h, E^h\right] = \sum_{s \in S^h} \phi_i^h(s) \mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^{hs}, E^{hs}\right], \tag{2}$$

with $\mathbb{E}[v(\tilde{\theta})|\theta^h, E^h] = v(\theta^h)$ at any $h$ such that the buyer chooses to take an action. We say that an experiment sequence $E$ is *fully informative* if $\mathbb{E}[v(\tilde{\theta})|\theta, E] = \theta\alpha + (1-\theta)\beta$; otherwise, it is *imperfect*. The fully informative one-shot experiment is a special case of fully informative experiment sequences. For brevity, we denote the value of a fully informative sequence by $VI(\theta)$. Clearly, $VI(\theta) \geq VI(\theta; E)$ for any $\theta$ and $E$. Throughout the analysis, we maintain the following assumption about the two types.

**Assumption 1.** *$VI(\theta_H) > VI(\theta_L)$ and $\theta_L > \theta_H$.*

The first part of this assumption is just a convention for identifying type $\theta_H$ as the type with a higher demand for fully informative experiment sequences; we will simply refer to type $\theta_H$ as the high type and type $\theta_L$ as the low type. The second part allows us to be more specific with respect to the direction of belief changes. Our exposition will focus on the case where the high type has a stronger belief about state $\omega_0$, and hence attaches a larger probability than does the low type that a given experiment will generate a signal that leads to downward belief revision. The opposite case (i.e., $VI(\theta_H) > VI(\theta_L)$ and $\theta_H < \theta_L$) is symmetric, but the exposition will be cluttered and repetitive if we try to deal with both cases under unified notation.

In our setup, the value of a fully informative sequence is $VI(\theta) = \min\{\theta\alpha, (1-\theta)\beta\}$, which is concave and reaches a maximum at $\theta = \sigma$. For $\theta_L > \theta_H$, there exists a $\underline{\theta} < \sigma$ such that Assumption 1 holds if and only if $\theta_L > \sigma$ and $\theta_H \in (\underline{\theta}, \theta_L)$. The assumption thus implies $\theta_L > \sigma$. Following the terminology of Bergemann et al. (2018), there are two cases. In the case of *congruent beliefs*, $\theta_L > \theta_H \geq \sigma$, the default actions of both types are $a_1$. In the case of *noncongruent beliefs*, $\theta_L > \sigma > \theta_H$, the default action of the low type is $a_1$ while that of the high type is $a_0$.

### 3.2. Seller's problem

Let $C(\theta_i; M)$ be the expected total payment of purchasing contract $M$ for type $i$. As usual, any optimal mechanism $(M_L, M_H)$ must satisfy the standard incentive compatibility and individual rationality constraints. These constraints require, for $i, i' = L, H$ and $i \neq i'$,

$$VI(\theta_i; E_i) - C(\theta_i; M_i) \geq VI(\theta_i; E_{i'}) - C(\theta_i; M_{i'}), \tag{IC-$i$}$$

$$VI(\theta_i; E_i) - C(\theta_i; M_i) \geq 0. \tag{IR-$i$}$$

**Lemma 1.** *In any optimal mechanism, (i) the experiment offered to the high type is fully informative; (ii) (IR-L) always binds; and (iii) at least one of (IC-H) and (IR-H) binds.*

The proofs of Lemma 1 and of other formal statements are in the appendix when they do not appear in the text. Parts (i) and (ii) are the familiar "no distortion at the top" and "no rent at the bottom" results (Mussa and Rosen, 1978; Bergemann et al., 2018). A special feature of our model is reflected in part (iii), which suggests that (IC-$H$) need not be binding.[6] Indeed, we later show that under some conditions, the seller can devise a mechanism in which (IR-$H$) binds along with (IR-$L$), and extract full surplus from both types.

On top of the IR and IC constraints, an optimal mechanism must also satisfy interim individual rationality constraints at each history, to make sure that the buyer follows through

[6] Although the high type has a greater willingness-to-pay for the fully informative experiment than does the low type, Assumption 1 does not imply that the former always has a greater willingness-to-pay for any experiment than the latter.

with paying for the experiments specified in the contract. Lemma 1 establishes that $E_H$ is fully informative. It is therefore without loss to focus on mechanisms that offer a static contract $M_H = (e_H, p_H)$, where $e_H$ is the fully informative one-shot experiment, so that interim individual rationality does not become an issue for the contract $M_H$. For the contract $M_L$, let $C(\theta_i^h; M_L^h)$ be the expected continuation payment at $h \in \mathscr{H}$ for type $i$. Interim individual rationality for type $i = L, R$ requires that at any history $h \in \mathscr{H}$,

$$VI(\theta_i^h; E_L^h) - C(\theta_i^h; M_L^h) \geq 0. \tag{IIR-$hi$}$$

It is clear that (IIR-$hi$) must hold for $i = L$, because the contract $M_L$ is chosen by the low type in a truth-telling equilibrium. If (IIR-$hi$) does not hold for $i = H$ at some history $h$, this would provide an opportunity for *double deviation* by the high type, meaning that the high type purchases the contract $M_L$ but does not comply with it when history $h$ obtains. The possibility of double deviation suggests that the incentive compatibility constraint for the high type may be more stringent than (IC-$H$), where the payoff from mimicking the low type assumes full compliance. It turns out, however, that double deviation will not be an issue: for any mechanism that violates (IIR-$hH$) at some history $h$, there is a weakly superior mechanism that satisfies (IIR-$hH$) at all histories $h$.

**Lemma 2.** *It is without loss of generality to focus on mechanisms that satisfy (IIR-hH) at all histories $h \in \mathscr{H}$.*

To see why no double deviation can be profitable, note that if a contract $M_L$ fails to satisfy some of the high type's IIR constraints, it must involve a contingency in which the high type would not participate (if he deviated and purchased $M_L$). The seller can lower a price in this contingency while raising the stage-1 price in a way to keep the low type's expected payment unchanged. This modification is profitable, as it unambiguously increases the high type's expected payment and hence reduces his information rent.

To summarize, since it is optimal to offer the high type the fully informative one-shot experiment at price $p_H$, the seller's revenue is $qp_H + (1-q)C(\theta_L; M_L)$, where

$$p_H = VI(\theta_H) - \max\{VI(\theta_H; E_L) - C(\theta_H; M_L), 0\}$$
$$C(\theta_L; M_L) = VI(\theta_L; E_L).$$

The high type's information rent is his payoff from mimicking the low type; in our model this payoff can be negative, in which case (IR-$H$) binds and the high type will obtain no surplus. It is useful to express the information rent differently by defining

$$\Delta C(M_L) := C(\theta_H; M_L) - C(\theta_L; M_L),$$
$$\Delta VI(E_L) := VI(\theta_H; M_L) - VI(\theta_L; M_L).$$

Using the fact that (IR-$L$) is binding, the information rent can be written as $\Delta VI(E_L) - \Delta C(M_L)$. We refer to $\Delta C(M_L)$ as the *payment gap*. It measures the difference in expected payments for contract $M_L$ between the two types. Such a payment gap emerges in a dynamic mechanism because different types have different beliefs about the likely outcomes of an experiment sequence. The quantity $\Delta VI(E_L)$ measures the difference in the willingness-to-pay for experiment sequence $E_L$ between the two types. The information rent can thus be reduced by two channels: distorting the quality of the experiment sequence to lower $\Delta VI(E_L)$, or adjusting the payment schedule to raise $\Delta C(M_L)$.

Substituting these definitions into the seller's revenue, the seller's problem is

$$\max_{M_L=(E_L,P_L)} q\left[VI(\theta_H) - \max\{\Delta VI(E_L) - \Delta C(M_L), 0\}\right] + (1-q)VI(\theta_L; E_L), \tag{3}$$

subject to (IIR-$hi$) for $i = L, H$ at all $h \in \mathscr{H}$. Note that we have dropped the constraint (IC-$L$). Our approach is to solve problem (3) subject to the IIR constraints only, and show that the solution to this relaxed problem satisfies (IC-$L$).

Optimal mechanisms may not be unique. Suppose an optimal mechanism $M_L = (E_L, P_L)$ satisfies $\Delta VI(E_L) - \Delta C(M_L) < 0$. Then, holding $E_L$ fixed, marginal changes in the payment schedule $P_L$ will not alter the value of the objective function (3), as long as the IIR constraints are not violated. In what follows, whenever an optimal mechanism admits multiple payment schedules that deliver the same revenue subject to the IIR constraints, we focus on the one that maximizes $\Delta C(M_L)$. Doing so is without loss of generality and it helps to highlight the role of the payment gap.

### 3.3. Characterization

The seller's objective function (3) depends both on the nature of the experiment sequence $E_L$ and on the structure of the payment schedule $P_L$. Although these two aspects are intertwined, we can still derive some useful properties of optimal mechanisms by examining each aspect separately. We begin with the properties of $E_L$.

The value of an experiment sequence depends only on the distribution of actions induced by the sequence in each state. We say that a history is a *terminal history* if no more information can be obtained afterwards. Any $h \in \mathscr{H}_{T+1}$ is a terminal history by definition. Also, for any $t < T+1$, $h \in \mathscr{H}_t$ is a terminal history if all subsequent experiments are uninformative. Once a terminal history is reached, the buyer takes an action and ends the experimentation process. Denote the set of terminal histories by $\overline{\mathscr{H}}$.

Due to the piecewise linearity of the value function, any terminal history $h \in \overline{\mathscr{H}}$ can be classified into one of three cases: (i) $\theta_L^h > \theta_H^h \geq \sigma$; (ii) $\theta_L^h \geq \sigma > \theta_H^h$; (iii) $\sigma > \theta_L^h > \theta_H^h$. Define

$$\Gamma_{ja_1} := \mathbb{P}\left[\theta_L^h > \theta_H^h \geq \sigma \mid h \in \overline{\mathscr{H}}, \omega = \omega_j\right],$$
$$\Gamma_{ja_0} := \mathbb{P}\left[\sigma > \theta_L^h > \theta_H^h \mid h \in \overline{\mathscr{H}}, \omega = \omega_j\right],$$

where the probabilities are induced by the experiment sequence $E_L$.

**Lemma 3.** *In any optimal mechanism, (i)* $\Gamma_{ja_0} + \Gamma_{ja_1} = 1$ *for* $j = 0, 1$*; and (ii)* $\Gamma_{1a_0} = 0$ *and* $\Gamma_{1a_1} = 1$.

*Proof.* (i) Suppose the statement is false. Then, there exists some terminal history $hs$ under $E_L$ such that $\theta_L^{hs} \geq \sigma > \theta_H^{hs}$. We can modify the stage experiment at history $h$ to $\hat{e}^h = (\hat{S}^h, \hat{\pi}^h)$, where $\hat{S}^h$ is obtained by deleting signal $s$ from the original $S^h$, and adding two signals, $s'$ and $s''$, to $\hat{S}^h$, which would split the belief $\theta_H^{hs}$ into $\hat{\theta}_H^{hs'} = 0$ and $\hat{\theta}_H^{hs''} = \sigma$. This operation does not affect $VI(\theta_H; E_L)$ as both 0 and $\sigma$ would induce the same action $a_0$, but the same operation would split $\theta_L^h$ into 0 and some $\hat{\theta}_L^{hs''} > \sigma$. As these two beliefs induce different actions, the operation is strictly beneficial to the low type. Thus, $VI(\theta_L; E_L)$ increases and $\Delta VI(E_L)$ decreases. The payment gap $\Delta C(M_L)$ is unaffected as long as only the stage-1 price $p_L$ in the original mechanism is adjusted. The overall effect is that the value of the objective function (3) increases. The original mechanism is not optimal.

(ii) Suppose $\Gamma_{1a_0} > 0$, so that there is a positive probability of having $h \in \overline{\mathscr{H}}$ such that $\sigma > \theta_L^h > \theta_H^h > 0$. We can then split this belief into 0 and 1. This modification increases $VI(\theta_L; E_L)$ by the amount $\theta_L \gamma_1^h \alpha > 0$, while it increases $VI(\theta_H; E_L)$ by a smaller amount $\theta_H \gamma_1^h \alpha$. Thus, $\Delta VI(E_L)$ decreases. The payment gap $C(M_L)$ is unaffected as long as only the stage-1 price $p_L$ in the original mechanism is adjusted. The total effect on the value of the objective function (3) is positive. This contradicts the optimality of the original mechanism. Hence $\Gamma_{1a_0} = 0$, which also implies $\Gamma_{1a_1} = 1$ by part (i). ∎

An implication of Lemma 3 is that the two types would always take the same action at all terminal histories. This feature makes the subsequent analysis more tractable. In particular, the difference in the value of $E_L$ can be written as

$$\Delta VI(E_L) = (\theta_L - \theta_H)(\Gamma_{0a_0}\beta - \alpha) - v(\theta_H) + v(\theta_L). \tag{4}$$

Another implication is that action $a_0$ can be chosen only when the state is known to be $\omega_0$, i.e., only when the belief is $\theta_i^h = 0$ at some terminal history $h$. An optimal experiment sequence can be either fully informative or imperfect; and if it is imperfect, it will entail action $a_1$ being chosen at belief $\theta_i^h < 1$ at some terminal history $h$. The reason why such a distortion can be profitable is the same as in the static case of Bergemann et al. (2018): lowering $\Gamma_{0a_0}$ from 1 reduces the value of the experiment sequence to the high type more than it does to the low type, thus reducing the information rent through reducing $\Delta VI(E_L)$. This property is preserved in our dynamic setting.

We now consider the problem of maximizing the payment gap. The expected payment for contract $M_L$ is given by

$$C(\theta_i; M_L) = p_L + \sum_{h \in \mathcal{H}} \left[\theta_i \gamma_1^h + (1-\theta_i)\gamma_0^h\right] p_L^h.$$

Therefore, the payment gap can be obtained as

$$\Delta C(M_L) = (\theta_L - \theta_H) \sum_{h \in \mathcal{H}} (\gamma_0^h - \gamma_1^h) p_L^h. \tag{5}$$

Equation (5) suggests two important features. First, since the stage-1 price $p_L$ is not contingent on the outcome of any experiment, changes in $p_L$ have no effect on the size of the payment gap. Second, it is optimal to set $p_L^h = 0$ at all $h$ such that $\gamma_0^h - \gamma_1^h \leq 0$, and adjust $p_L$ accordingly to maintain (IR-$L$). For the same reason, it is optimal to set at least a small positive price $p_L^h$ at histories $h$ such that $\gamma_0^h - \gamma_1^h > 0$. This shows that $\Delta C(M_L) > 0$ in an optimal mechanism.

For the subsequent analysis, it is convenient to define up signals and down signals. A signal $s \in S^h$ is an *up signal* if its likelihood ratio $\pi^h(s|\omega_1)/\pi^h(s|\omega_0)$ is greater than or equal to 1, and is a *down signal* otherwise. Bayes' rule implies that $\theta_i^{hs} \geq \theta_i^h$ for $i = L, H$ if $s$ is an up signal, and $\theta_i^{hs} < \theta_i^h$ if $s$ is a down signal.

**Lemma 4.** *It is without loss of generality to assume that in an optimal mechanism, any history followed by an up signal is a terminal history leading to action $a_1$.*

In the proof of Lemma 4, we show that there is a way to lower $p_L^{hs}$ and raise $p_L^h$ to increase the payment gap whenever $s$ is an up signal, because $\gamma_1^{hs}/\gamma_0^{hs} \geq \gamma_1^h/\gamma_0^h$. Thus it is optimal to set $p_L^{hs} = 0$ at such histories.[7] Given that any history followed by an up signal does not contribute positively to the seller's revenue or to the payment gap, there is an equivalent mechanism that will remove this history from $\mathscr{H}$ without affecting the seller's revenue or the IIR constraints, unless this history is a terminal history.

Suppose $s', s'' \in S^h$ are both up signals. By Lemma 4, $hs'$ and $hs''$ are both terminal histories that would result in the same action (namely, action $a_1$) for both types. These two histories can be merged into a single terminal history, denoted generically as $hu$, that leads to action $a_1$ with the same combined probability without affecting the value of the experiment sequence or any of the constraints. Thus, another implication of Lemma 4 is that it is without loss to consider experiment sequences in which at every history $h$, there is only one element $u \in S^h$ such that $\pi^h(u|\omega_1)/\pi^h(u|\omega_0) \geq 1$.

While Lemma 4 shows the buyer takes an action immediately after an up signal, the next result shows that he never takes an action immediately after a down signal unless he is sure that the state is $\omega_0$.

**Lemma 5.** *In an optimal mechanism, the experiment sequence continues until an up signal is observed or until the belief reaches 0.*

Intuitively, if an experiment $e^h$ with down signal $s$ would lead to beliefs such that $\theta_L^{hs} > \theta_H^{hs} \geq \sigma$, then the demand for information by both types at history $hs$ could be increased by lowering their beliefs closer toward $\sigma$. The proof of Lemma 5 shows that in an optimal mechanism, a down signal $s$ must result in a belief $\theta_H^{hs} < \sigma$. Therefore, if the buyer takes an action immediately after a down signal, he will choose to take action $a_0$. But then this implies that the buyer can only take action $a_0$ at history $hs$ only when $\theta_L^{hs} = \theta_H^{hs} = 0$, for otherwise it would violate Lemma 3.

To further analyze the problem of maximizing the payment gap, it is useful to characterize the optimal payment schedule. Lemma 6 below allows us to express the price charged at each history $h$ as the value of the stage-experiment $e_L^h$ for the high type, *as if* the high type were to take a final action after this experiment.

**Lemma 6.** *In an optimal mechanism, (i) (IIR-hH) binds at any $h \in \mathscr{H}$; and (ii) $p_L^h = VI(\theta_H^h; e_L^h)$ at any $h \in \mathscr{H}$.*

This result allows us to reduce the problem of choosing $M_L = (E_L, P_L)$ to a problem of choosing $E_L$ alone. Once $E_L$ is chosen, the prices $p_L^h$ are determined by part (ii) of the

[7] This argument generalizes our earlier point that it is optimal to set $p_L^h = 0$ whenever $\gamma_1^h/\gamma_0^h \geq 1$.

lemma for any $h \in \mathscr{H}$. The stage-1 price $p_L$ is in turn determined by (IR-$L$), which gives

$$p_L = VI(\theta_L; E_L) - \sum_{h \in \mathscr{H}} \left(\theta_L \gamma_1^h + (1-\theta_L)\gamma_0^h\right) p_L^h.$$

Moreover, part (ii) of Lemma 6 implies part (i). Once $\{p_L^h\}_{h \in \mathscr{H}}$ is substituted into equation (5) for the payment gap $\Delta C(M_L)$, the only remaining constraints for revenue maximization problem (3) are the IIR constraints for the low type.

### 3.4. Optimal binary mechanism

So far, we have only shown that it is without loss to consider mechanisms in which the signal space comprises one up signal and possibly multiple down signals. An even simpler class of mechanisms is *binary mechanisms*, in which all experiments in the experiment sequence are binary, yielding either an up signal or a down signal. By Lemma 4, in a binary mechanism, the buyer purchases an additional experiment only after a down signal. Given this, with slight abuse of notation, the relevant history in a binary mechanism can be represented by $t \in \{1, \dots, T+1\}$. For $t > 1$, history $t$ indicates a history where the buyer has observed down signals at all stages from 1 to $t-1$; for $t = 1$, history $t$ indicates the initial history. The experiment at history $t$ can be written as $e_L^t = (\{u, d\}, \pi^t)$, where a binary experiment $e_L^t$ is represented by two numbers:

$$\pi_{0d}^t := \pi^t(d \mid \omega_0), \quad \pi_{1u}^t := \pi^t(u \mid \omega_1).$$

The interpretation of $u$ as an up signal requires the convention that $\pi_{1u}^t \geq 1 - \pi_{0d}^t$. If $e^t$ is fully informative, then $\pi_{1u}^t = \pi_{0d}^t = 1$; and if $e^t$ is totally uninformative, then $\pi_{1u}^t = 1 - \pi_{0d}^t$. The following proposition establishes that it is without loss of generality to focus on binary mechanisms and provides some important properties of optimal binary mechanisms.

**Proposition 1.** *An optimal mechanism can be achieved by a binary mechanism. In an optimal binary mechanism, (i)* $\pi_{0d}^t = 1$ *at* $t = 2, \dots, T$*; and (ii)* $\pi_{1u}^T = 1$.

Because $\pi_{1u}^T = \pi_{0d}^T = 1$, Proposition 1 implies that the stage-$T$ experiment $e^T$ must be fully informative. Moreover, because $\pi_{0d}^t = 1$ for all $t \geq 2$, the proposition also implies that from stage 2 onward, the continuation experiment sequence $E_L^t$ must be fully informative, eventually leading to the correct action chosen in each state. Any distortion of an optimal experiment sequence $E$ can occur only if $\pi_{0d}^1 < 1$, where an up signal in state $\omega_0$ causes the buyer to incorrectly choose $a_1$ at the first stage. In this sense, distortions, if any, are

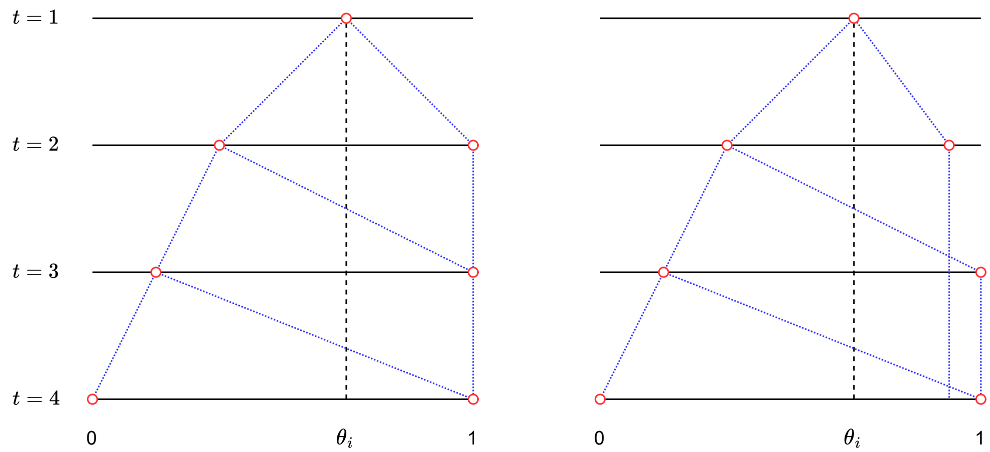


***Figure 1.*** *The optimal binary mechanism must induce beliefs with one of the two structures. In the left panel, the experiment sequence is fully informative. In the right panel, the experiment sequence is imperfect, with an imperfect experiment offered at stage 1.*

*front-loaded* in an optimal mechanism. Proposition 1 leads to a simple characterization of how the belief evolves under an optimal binary experiment sequence $E_L$. If $E_L$ is fully informative, then an up signal at every stage will lead to a belief of 1 and action $a_1$ being chosen, while a down signal in each stage will cause the belief to gradually go down until it reaches 0 at the last stage. If $E_L$ is imperfect, an up signal in stage 1 will lead to a belief strictly below 1, while up signals in all other stages will lead to a belief of 1. See Figure 1 for an illustration.

We prove Proposition 1 using a "front-loading" argument. Suppose $\pi^t_{0d} < 1$ for some $t \geq 2$. We could then (i) increase the distortion at stage $t-1$ by lowering $\pi^{t-1}_{0d}$ (thereby lowering the value of the stage experiment $e^{t-1}_L$ and hence lowering $p^{t-1}_L$), (ii) reduce the distortion at stage $t$ by raising $\pi^t_{0d}$ (thereby raising $p^t_L$), and (iii) adjust $\pi^t_{1u}$ and $\pi^{t-1}_{1u}$ in a way to strictly increase the payment gap $\Delta C(M_L)$. Intuitively, a positive payment gap arises because the ratio of probabilities (held by the high type to the low type) of observing $t-1$ consecutive down signals is larger than 1, but the corresponding ratio of probabilities of observing $t$ consecutive down signals is even higher, making it profitable to front-load distortions (and accordingly back-load prices). We establish this front-loading argument without assuming that experiments are binary, but Proposition 1 also shows that any optimal mechanism can be reduced to a binary mechanism.

Proposition 1 substantially simplifies the seller's problem, as she now only needs to determine $\{\pi_{1u}^t\}_{t=1}^{T-1}$ and $\pi_{0d}^1$. Alternatively, once $\pi_{0d}^1$ is fixed, the problem can be equivalently formulated as choosing a belief sequence $\{\theta_H^t\}_{t=2}^T$, where $\theta_i^t$ denotes type $i$'s belief at history $t$ (following $t-1$ consecutive down signals). Observe that once $\theta_H^t$ is determined, and given the initial belief $\theta_H$, the corresponding $\theta_L^t$ is uniquely pinned down. Let $\rho(\theta_H^t;\theta_H)=\theta_L^t$ denote this association.

Since for $t=2,\ldots,T$,

$$\gamma_0^t=\pi_{0d}^1,\quad \gamma_1^t=\frac{\theta_H^t}{1-\theta_H^t}\frac{1-\theta_H}{\theta_H}\pi_{0d}^1,\quad p_L^t=\theta_H^t\pi_{1u}^t\alpha=\frac{\theta_H^t-\theta_H^{t+1}}{1-\theta_H^{t+1}}\alpha,$$

the payment gap (5) can be written as

$$\Delta C=\frac{\theta_L-\theta_H}{\theta_H}\pi_{0d}^1\alpha\sum_{t=2}^{T}\frac{(\theta_H-\theta_H^t)(\theta_H^t-\theta_H^{t+1})}{(1-\theta_H^t)(1-\theta_H^{t+1})}. \tag{6}$$

This implies that the optimal belief sequence that maximizes the payment gap is determined independently of $\pi_{0d}^1$. We can therefore solve the problem in two steps: first derive the optimal belief sequence that maximizes the payment gap, and then determine the distortion level $\pi_{0d}^1$ taking the belief sequence as given.

The payment gap (6) is decreasing in $\theta_H^{T+1}$; so it is optimal to set $\theta_H^{T+1}=0$. For $t=3,\ldots,T$, the derivative of $\Delta C$ with respect to $\theta_H^t$ is proportional to

$$-\frac{(1-\theta_H)(\theta_H^t-\theta_H^{t+1})}{(1-\theta_H^t)^2(1-\theta_H^{t+1})}-\frac{\theta_H-\theta_H^{t-1}}{(1-\theta_H^t)^2}+\frac{\theta_H-\theta_H^t}{(1-\theta_H^t)(1-\theta_H^{t+1})}.$$

The seller can increase $\theta_H^t$ by reducing $\pi_{1u}^{t-1}$. Holding everything else constant, increasing $\theta_H^t$ has two effects. First, it reduces the gap in the conditional probability of reaching history $t$ between the two types, which is captured by the first term. Second, a higher $\theta_H^t$ makes the stage-experiment $e_L^{t-1}$ less informative and the stage-experiment $e_L^t$ more informative, thus reducing $p_L^{t-1}$ and raising $p_L^t$, as captured by the second and third terms. Setting this derivative to 0 gives rise to a second-order difference equation,

$$(\theta_H^{t-1}-\theta_H^t)(1-\theta_H^{t+1})-(\theta_H^t-\theta_H^{t+1})(1-\theta_H^t)=0,$$

with boundary conditions $\theta_H^2$ and $\theta_H^{T+1}=0$. The explicit solution is

$$\theta_H^t=\psi^t(\theta_H^2):=1-(1-\theta_H^2)^{\frac{T-t+1}{T-1}},$$

for $t = 2, \dots, T+1$. This suggests that the entire belief sequence is uniquely determined once we pin down $\theta_H^2$. Define

$$G_T(\theta_H^2; \theta_H) := \frac{\alpha}{\theta_H} \sum_{t=2}^{T} \frac{(\theta_H - \psi^t(\theta_H^2))(\psi^t(\theta_H^2) - \psi^{t+1}(\theta_H^2))}{(1-\psi^t(\theta_H^2))(1-\psi^{t+1}(\theta_H^2))}. \tag{7}$$

Then the payment gap (6) can be written as

$$\Delta C = (\theta_L - \theta_H)\pi_{0d}^1 G_T(\theta_H^2; \theta_H). \tag{8}$$

The problem of selecting $\theta_H^2$ is qualitatively different from that of selecting other beliefs because the solution also needs to satisfy (IIR-2$L$).[8] To proceed, let

$$\Delta C^2(\theta_H^2; \theta_H) := \left(\rho(\theta_H^2; \theta_H) - \theta_H^2\right) \pi_{0d}^1 \left[ \frac{\alpha}{\theta_H^2} \sum_{t=3}^{T} \frac{(\theta_H^2 - \psi^t(\theta_H^2))(\psi^t(\theta_H^2) - \psi^{t+1}(\theta_H^2))}{(1-\psi^t(\theta_H^2))(1-\psi^{t+1}(\theta_H^2))} \right]$$

denote the payment gap in the continuation mechanism at stage 2 after the high type's belief has become $\theta_H^2$. Since the continuation mechanism is fully informative by Proposition 1, (IIR-2$L$) can be written as

$$VI(\rho(\theta_H^2; \theta_H)) - VI(\theta_H^2) + \Delta C^2(\theta_H^2; \theta_H) \geq 0. \tag{9}$$

Let $\theta_T^*$ be the value of $\theta_H^2$ for which (9) holds with equality; and if (9) holds for all $\theta_H^2 \leq \theta_H$ let $\theta_T^* = \theta_H$. Also let $\theta_T^{**} = \arg\max_{\theta_H^2} G_T(\theta_H^2; \theta_H)$.

**Lemma 7.** *Maximizing $\Delta C$ subject to (IIR-2$L$) requires setting $\theta_H^2 = \min\{\theta_T^*, \theta_T^{**}\}$, where $\theta_T^{**} = 1 - (1-\theta_H)^{\frac{T-1}{T}}$ and $\theta_T^* \in (\underline{\theta}, \sigma)$ is uniquely defined.*

Define

$$G_T^* := G_T(\min\{\theta_T^*, \theta_T^{**}\}; \theta_H).$$

Since $\Gamma_{0a_0} = \pi_{0d}^1$ under an optimal experiment sequence $E_L$, using equations (4) and (8), the high type's information rent is minimized at

$$\Delta VI - \Delta C = (\theta_L - \theta_H)(\beta - G_T^*)\pi_{0d}^1 - (v(\theta_H) - \theta_H \alpha).$$

When $\Delta VI - \Delta C < 0$ at $\pi_{0d}^1 = 1$, there is no tradeoff between minimizing the information rent and maximizing the value of $E_L$. Otherwise, raising $\pi_{0d}^1$ would increase $VI(\theta_L; E_L)$ to the low type and increase the information rent to the high type; the seller's revenue in this case is linear in $\pi_{0d}^1$ and the tradeoff depends the magnitude of $q$ relative to $1-q$. We thus arrive at the following characterization of an optimal mechanism.

[8] This problem arises only at $t = 2$ because (IIR-2$L$) implies (IIR-$tL$) for $t \geq 3$, as will be shown.

**Proposition 2.** *In an optimal binary mechanism, the high type is offered the static fully informative experiment. The low type is offered experiment sequence $E_L$, which would induce beliefs $\{\theta_H^t\}_{t=2}^{T+1}$ on the high type upon $t-1$ consecutive down signals, with*

$$\theta_H^2 = \min\{\theta_T^*, \theta_T^{**}\}, \quad \theta_H^t = 1-(1-\theta_H^2)^{\frac{T-t+1}{T-1}} \text{ at } t = 2,\dots,T+1.$$

*Moreover,*

1. *If $v(\theta_H)-\theta_H\alpha \geq (\theta_L-\theta_H)(\beta-G_T^*)$, then $\pi_{0d}^1 = 1$. Experiment $E_L$ is fully informative, and there is full surplus extraction.*
2. *If $v(\theta_H)-\theta_H\alpha < (\theta_L-\theta_H)(\beta-G_T^*)$, there are two cases.*
   - *(a) If $(1-q)(1-\theta_L)\beta \geq q(\theta_L-\theta_H)(\beta-G_T^*)$, then $\pi_{0d}^1 = 1$. Experiment $E_L$ is fully informative, but the high type retains some positive surplus.*
   - *(b) If $(1-q)(1-\theta_L)\beta < q(\theta_L-\theta_H)(\beta-G_T^*)$, then*

$$\pi_{0d}^1 = \frac{v(\theta_H)-\theta_H\alpha}{(\theta_L-\theta_H)(\beta-G_T^*)}.$$

*Experiment $E_L$ is imperfect, and neither type obtains positive surplus.*

The values of $\pi_{0d}^1$ and $\theta_H^2$ described in Proposition 2 together determine $\pi_{1u}^1$, which in turn pin down the high type's belief upon observing an up signal in experiment $e_L^1$. This belief is equal to 1 if $\pi_{0d}^1 = 1$, or strictly less than 1 if $\pi_{0d}^1 < 1$. The belief upon observing an up signal in experiment $e_L^t$ at $t = 2,\dots,T$ is equal to 1, as implied by Proposition 1. Propositions 1 and 2 therefore provide a complete characterization of the beliefs induced by the optimal experiment sequence.

Proposition 2 is obtained by utilizing all the previously established results to reduce the objective function of problem (3) into a function of $\pi_{0d}^1$ only, and then maximizing over $\pi_{0d}^1$. We complete the proof by showing that the mechanism so derived satisfies (IIR-$ti$) for $i = L, H$ at all $t = 2,\dots,T$. Finally, it is clear that (IC-$L$) holds, because the low type has no incentive to choose $M_H$ under Assumption 1.

Although our model setting considers $T \geq 2$, Proposition 2 formally applies to static mechanisms with $T = 1$ as well. The payment gap is equal to 0 in the static case, which implies that $G_1^* = 0$. For $G_1^* = 0$, Cases 2(a) and 2(b) of the proposition correspond to the optimal static mechanism described in Bergemann et al. (2018), while Case 1 never obtains. This suggests that full surplus extraction is infeasible if mechanisms are restricted to be static.

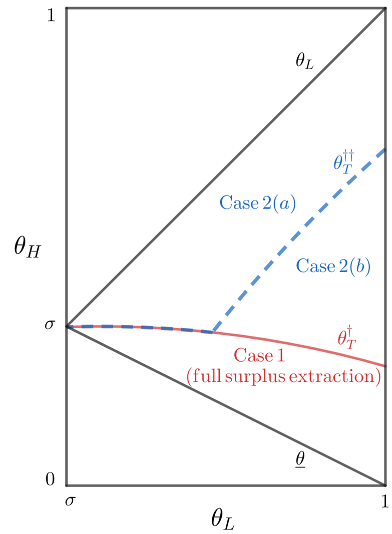


***Figure 2.*** *The parameter space is divided into three regions according to Proposition 3.*

# 4. Discussion

## 4.1. Full surplus extraction and the gains from a dynamic mechanism

While Proposition 2 establishes a necessary and sufficient condition for an optimal (binary) mechanism, it involves an endogenous object $G_T^*$ and is not stated in terms of model primitives. Here, we provide a characterization of the condition for which each of the three cases in Proposition 2 obtains. A key parameter in this regard is $\theta_H$, which measures the dispersion in initial beliefs for a given $\theta_L$.

**Proposition 3.** *For any $T \geq 2$, there is a unique $\theta_T^{\dagger} \in (\underline{\theta}, \sigma)$ and a unique $\theta_T^{\dagger\dagger} \in [\theta_T^{\dagger}, \theta_L)$ such that (i) if $\theta_H \leq \theta_T^{\dagger}$, Case 1 of Proposition 2 obtains; (ii) if $\theta_H \in (\theta_T^{\dagger}, \theta_T^{\dagger\dagger})$, Case 2(b) obtains; and (iii) if $\theta_H \geq \theta_T^{\dagger\dagger}$, Case 2(a) obtains. Moreover, $\theta_T^{\dagger}$ decreases in $\theta_H$.*

Figure 2 illustrates Proposition 3. For a given $\theta_L > \sigma$, Assumption 1 requires $\theta_H \in (\underline{\theta}, \theta_L)$. Figure 2 shows that the region where full surplus extraction obtains (i.e., $\theta_H \in (\underline{\theta}, \theta_T^{\dagger}]$) is always nonempty for any given $\theta_L$. Full surplus extraction is possible when there is a large dispersion in initial beliefs—$\theta_H$ is sufficiently far from $\theta_L$—because a larger dispersion leads to a larger payment gap. Indeed, in the case of congruent beliefs (i.e., $\theta_L > \theta_H \geq \sigma$), the seller cannot extract full surplus.

Even if full surplus extraction is not feasible for some $T$, it is straightforward to show that the seller can increase her revenue by introducing a dynamic mechanism with one more stage, as long as there still remains surplus to be extracted from trade. Since no static mechanisms can achieve full surplus extraction, this fact implies that adding a second stage to a static mechanism is always strictly beneficial for the seller.

**Proposition 4.** *The seller strictly benefits from having an additional stage unless the mechanism has already achieved full surplus extraction. Moreover, $\theta_T^\dagger$ strictly increases in $T$.*

*Proof.* Fix an optimal mechanism for some $T$. If there is an additional stage, we can add $\theta_H^t - \varepsilon$ to the belief sequence, where $\varepsilon > 0$ is some arbitrarily small number that satisfies $\theta_H^t > \theta_H^t - \varepsilon > \theta_H^{t+1}$ for some $t = 2, \ldots, T$. This operation raises the payment gap if

$$\frac{(\theta_H - \theta_H^t)(\theta_H^t - \theta_H^t + \varepsilon)}{(1-\theta_H^t)(1-\theta_H^t + \varepsilon)} + \frac{(\theta_H - \theta_H^t + \varepsilon)(\theta_H^t - \theta_H^{t+1} - \varepsilon)}{(1-\theta_H^t + \varepsilon)(1-\theta_H^{t+1})} > \frac{(\theta_H - \theta_H^t)(\theta_H^t - \theta_H^{t+1})}{(1-\theta_H^t)(1-\theta_H^{t+1})}.$$

This condition reduces to $(\theta_H^t - \theta_H^{t+1} - \varepsilon)\varepsilon > 0$, which holds with our choice of $\varepsilon$.

This argument also shows that $G_T^*$ increases in $T$. Since $\theta_T^\dagger$ is determined by the condition that $v(\theta_H) - \theta_H \alpha - (\theta_L - \theta_H)(\beta - G_T^*) = 0$, and the left-hand side of this equation is decreasing in $\theta_H$, it follows that $\theta_T^\dagger$ increases in $T$. ∎

### 4.2. Properties of optimal stage prices and stage experiments

Our characterization of optimal dynamic mechanisms has some interesting implications for the structure of the stage prices and stage experiments. In particular, observe that the stage-1 price $p_L$ in the optimal price schedule $P_L$ satisfies

$$p_L = C(\theta_L; M_L) - \phi_L^1(d) C(\theta_L^2; M_L^2) \geq VI(\theta_L; E_L) - \phi_L^1(d) VI(\theta_L^2) = VI(\theta_L; e_L^1) \geq 0. \quad (10)$$

This implies that $p_L = 0$ if and only if the following conditions hold: (i) (IIR-2$L$) binds (so that the first inequality becomes an equality); and (ii) the value of the first-stage experiment $e_L^1$ to the low type is 0 (so that the second inequality becomes an equality).

In an online appendix to their paper, Bergemann et al. (2018) consider a two-stage mechanism with congruent beliefs and show that it can improve the seller's revenue by means of an example. The example they construct has two properties: (i) the stage-1 experiment is provided for free; and (ii) the stage-1 experiment yields a down signal that equates the values of information for the two types, i.e., $\theta_H^2$ satisfies $\theta_H^2 \alpha = (1 - \rho(\theta_H^2; \theta_H))\beta$. Equation (10) shows that the seller can charge a positive price for the stage-1 experiment whenever (IIR-2$L$) is slack. Moreover, Lemma 7 shows that the choice of $\theta_H^2$ in their example is not always optimal, as maximizing the payment gap subject to (IIR-2$L$) requires setting $\theta_H^2 = \min\{\theta_T^*, \theta_T^{**}\}$.[9]

**Proposition 5.** *In an optimal binary mechanism, (i) the prices are constant over stages, i.e., $p_L^{t+1} = p_L^t$ at $t = 2, \ldots, T-1$; and (ii) the stage-experiment $e_L^{t+1}$ is more Blackwell-informative than the stage-experiment $e_L^t$ at $t = 2, \ldots, T-1$.*

*Proof.* Lemma 6 requires each stage price to be equal to the value of the stage experiment. Therefore,

$$p_L^t = \theta_H^t \pi_{1u}^t \alpha = \frac{\theta_H^t - \theta_H^{t+1}}{1 - \theta_H^{t+1}} \alpha.$$

Using the explicit formula, $\theta_H^t = 1 - (1 - \theta_H^2)^{\frac{T-t+1}{T-1}}$ for $t = 2, \ldots, T$, we obtain

$$p_L^t = \left[1 - (1 - \theta_H^2)^{\frac{1}{T-1}}\right] \alpha,$$

which is independent of $t$.

Recall also that $\pi_{0d}^t = 1$ at all $t = 2, \ldots, T$. Given this, the experiment $e_L^{t+1}$ is more Blackwell-informative than $e_L^t$ if

$$\pi_{1u}^{t+1} = \frac{\theta_H^{t+1} - \theta_H^{t+2}}{\theta_H^{t+1}(1 - \theta_H^{t+2})} > \frac{\theta_H^t - \theta_H^{t+1}}{\theta_H^t(1 - \theta_H^{t+1})} = \pi_{1u}^t.$$

Since $p_L^t$ is constant, this condition reduces to $\theta_H^t > \theta_H^{t+1}$. ∎

A remarkable feature of the optimal mechanism is that the seller offers a constant price across stages. Since the price charged at each stage is simply the value of the stage experiment, this alternatively implies that the optimal belief sequence is the one that equates the values of the stage experiments across stages. Observe that the buyer gains more information, and hence the demand for information declines, as he observes more and more down signals over stages. The fact that the price remains constant then implies that the seller must provide more precise information at later stages, giving rise to the second statement that experiments become progressively more informative in the sense of Blackwell.

[9] Choosing $\theta_H^2$ satisfying $\theta_H^2 \alpha = (1 - \rho(\theta_H^2; \theta_H))\beta$ is optimal if $T = 2$ and (IIR-2$L$) is binding.

### 4.3. Divergence in beliefs and surplus extraction

Hörner and Skrzypacz (2016) consider a persuasion game between an informed seller and an uninformed buyer and show that the seller, with no commitment power, can improve her revenue by disclosing information gradually over time. While we consider a totally different economic problem—a mechanism-design problem in which an uninformed seller screens partially informed buyers—there is an important commonality regarding the way surplus extraction occurs along a dynamic dimension. In the context where some concerned parties have private information, they necessarily hold different expectations about how the belief evolves even when they face the same information-generating process. Our analysis reveals that a party who controls the flow of information can extract more surplus by leveraging this divergence in beliefs.

In Hörner and Skrzypacz (2016), the party who controls the flow of information is the (informed) competent type, as the incompetent type can only mimic the competent type in equilibrium. At any point, the buyer's willingness to pay is bounded by his current belief, which is always below the competent type's valuation. As a consequence, the competent type can never appropriate the full value of her private information. However, since the competent type knows that she is more likely to pass any given pass-or-fail test than the uninformed buyer, there emerges a divergence in perceived payments between the two parties. This divergence allows the competent type to benefit from deferring payments, by withholding some information for later rounds.

In our setting, it is the seller who has full control over the flow of information. Given the same experiment sequence, the high type expects to reach the next round (by observing a down signal) with higher probability than the low type. Consequently, when payments are back-loaded, the expected payment is always higher for the high type than for the low type. This discourages the high type from mimicking the low type, thereby reducing the high type's information rent without distorting the quality of information and even enabling the seller to extract full surplus under some conditions. Moreover, since this channel can substitute for the conventional channel of quality distortion, the optimal dynamic mechanism in general entails less quality distortion than the optimal static mechanism: whenever the optimal dynamic mechanism provides imperfect information, the optimal static mechanism necessarily does so as well.

This perspective also sheds light on the *direction* of information disclosure. Our analysis shows that the deciding factor is whether the high type—the one with greater demand for information—has a higher or lower initial belief than the low type. We consider a setting in which the high type holds a lower initial belief. In this case, the optimal mechanism gradually collects payments after down signals until the belief reaches 0. If we instead consider a setting in which, contrary to Assumption 1, $VI(\theta_H) > VI(\theta_L)$ and $\theta_H > \theta_L$, then the direction reverses: the mechanism gradually collects payments after up signals until the belief reaches 1.

### 4.4. Infinite horizon

We have thus far taken the number of stages $T$ as exogenously given. As the mechanism necessarily becomes more complicated as $T$ increases, one can interpret this as the maximum complexity of mechanisms that the seller can offer. Proposition 4 shows that the seller benefits strictly from increasing the number of stages, as long as full surplus extraction is not yet achieved. It is useful to consider the limiting case as the number of stages goes to infinity, in order to obtain a theoretical upper bound of what the seller can achieve in our environment. Equivalently, one may also think of the limiting case as a model with fixed horizon, where the seller can provide experiments in continuous time.

When $T$ is infinitely large, the seller can provide information by an infinitesimally small margin, implying $\theta_H^t - \theta_H^{t+1} \to 0$ for all $t \geq 2$, while $\theta_H^2$ may still be bounded away from $\theta_H$. Observe that from equation (7),

$$G_\infty(\theta_H^2; \theta_H) := \lim_{T\to\infty} G_T(\theta_H^2; \theta_H) = \frac{\alpha}{\theta_H} \lim_{T\to\infty} \sum_{t=2}^{T} (\theta_H - \theta_H^t) \left( \frac{1}{1-\theta_H^t} - \frac{1}{1-\theta_H^{t+1}} \right),$$

with $\theta_H^t = \psi^t(\theta_H^2)$ for $t > 2$. Let $\theta_H(s)$ be the continuous limit of the belief sequence with $\theta_H(0) = \theta_H^2$ and $\theta_H(1) = 0$. In the limit, we obtain

$$G_\infty(\theta_H^2; \theta_H) = \frac{\alpha}{\theta_H} \int_0^1 (\theta_H(s) - \theta_H) \,\mathrm{d}\frac{1}{1-\theta_H(s)} = -\frac{\alpha}{\theta_H}(\theta_H + \ln(1-\theta_H^2)).$$

For the next result, define $\theta^\dagger \in (\underline{\theta}, \sigma)$ and $\theta^*$ such that they satisfy:

$$(1-\theta^\dagger)\beta - \theta^\dagger\alpha - (\theta_L - \theta^\dagger)\big(\beta - G_\infty(\theta^\dagger; \theta^\dagger)\big) = 0, \tag{11}$$

$$(1-\theta^*)\beta - \theta^*\alpha - (\rho(\theta^*; \theta_H) - \theta^*)(\beta - G_\infty(\theta^*; \theta^*)) = 0. \tag{12}$$

Note that the value of $\theta^*$ depends on $\theta_H$, while the value of $\theta^\dagger$ does not.

**Proposition 6.** *As the number of stages grows,*

$$\lim_{T\to\infty} \theta_T^\dagger = \theta^\dagger, \quad \lim_{T\to\infty} \theta_T^{**} = \theta_H, \quad \lim_{T\to\infty} \theta_T^* = \theta^* \begin{cases} > \theta_H & \text{if } \theta_H < \theta^\dagger \\ \leq \theta_H & \text{if } \theta_H \geq \theta^\dagger. \end{cases}$$

*In a model where the seller can choose any finite number of stages, the optimal dynamic mechanism extracts full surplus if and only if $\theta_H < \theta^\dagger$.*

When $\theta_H > \theta^\dagger$, no mechanisms can achieve full surplus extraction, no matter how complicated they can be. In this case, in the limit, the belief jumps down to $\theta^*$ after a down signal at stage 1 and then continuously declines over stages after that. Otherwise, full surplus extraction is feasible, and the belief continuously declines from the beginning with no discrete jump. We can therefore conclude that $\theta_H$ and $\theta_H^2$ are bounded away from each other if and only if full surplus extraction is not feasible.

Since $\theta_T^\dagger$ increases monotonically with $T$ and reaches a limit of $\theta^\dagger$, there exists $T'$ sufficiently large such that $\theta_H \leq \theta_{T'}^\dagger$ for any $\theta_H < \theta^\dagger$. By Proposition 3, the seller can extract full surplus by choosing a finite dynamic mechanism with $T'$ stages.

### 4.5. General value function

Although we develop the full characterization of optimal dynamic mechanisms under a setting with binary actions in binary states, the main insights of our model extend more broadly. To show this, we now consider the case where $v(\cdot)$ is an arbitrary convex value function. This extension allows for a general (finite or infinite) action space of the buyer in a binary state model.

With arbitrary value functions, even the optimal static mechanism is difficult to characterize as the optimal experiment may be more complex than a binary experiment. Haghpanah and Siegel (2025) make some progress in studying static screening problems where the single-crossing property need not hold, but extending their method to dynamic mechanisms is still an intractable open problem. In this section, we confine ourselves to making two observations.

**Proposition 7.** *For any convex value function $v(\cdot)$ and any $T \geq 2$, (i) there exists $\theta_T^\dagger > \underline{\theta}$ such that full surplus extraction is feasible if $\theta_H \leq \theta_T^\dagger$; and (ii) an optimal mechanism with $T$ stages yields strictly higher revenue than one with $T-1$ stages unless the latter has already achieved full surplus extraction.*

*Proof.* (i) Observe that a necessary and sufficient condition for full surplus extraction is

$$VI(\theta_H) - VI(\theta_L) \leq \Delta C(M_L).$$

As $\theta_H$ approaches $\underline{\theta}$ from above, the left-hand side converges to 0 while the right-hand side is bounded away from 0. To see the latter, consider a two-stage binary mechanism in which a down signal $d$ in the initial stage leads to a belief $\theta_H^2 \in (0, \underline{\theta})$. The seller charges $p_L^2 = VI(\theta_H^2) = \theta_H^2 v(1) + (1-\theta_H^2)v(0) - v(\theta_H^2) > 0$ for the fully informative experiment at the second stage after a down signal. Since $\pi(d|\omega_0) - \pi(d|\omega_1) > 0$ by definition the payment gap is

$$\Delta C(M_L) = (\theta_L - \theta_H)(\pi(d|\omega_0) - \pi(d|\omega_1))p_L^2,$$

which is positive and bounded away from 0. This proves that full surplus extraction is feasible if $\theta_H$ is sufficiently close to $\underline{\theta}$.

(ii) Fix an optimal mechanism with $T-1$ stages. Pick some history $h \in \mathscr{H}_{T-1}$. Let $e_L^h = (S^h, \pi^h)$ be the experiment offered at this history with $S^h = \{s_1, \ldots, s_N\}$, where $s_1$ is the signal that leads to the smallest belief. This means that $\pi^h(s_1|\omega_0) > \pi^h(s_1|\omega_1)$. We now construct a $T$-stage mechanism which is identical to the original mechanism except at history $h$. Specifically, at this history, the continuation experiment sequence $\hat{E}^h$ consists of two stages. The seller first offers an experiment $\hat{e}^1 = (S^h, \hat{\pi}^1)$ such that, for $j = 0, 1$,

$$\hat{\pi}^1(s|\omega_j) = \begin{cases} \varepsilon\pi^h(s|\omega_j) + 1 - \varepsilon & \text{if } s = s_1 \\ \varepsilon\pi^h(s|\omega_j) & \text{if } s \neq s_1, \end{cases}$$

for some $\varepsilon \in (0,1)$. At stage $T$, if the history is $hs_1$, the seller offers another experiment $\hat{e}^2 = (S^h, \hat{\pi}^2)$ such that

$$\hat{\pi}^2(s|\omega_j) = \begin{cases} \frac{1}{\hat{\pi}^1(s_1|\omega_j)}\pi^h(s|\omega_j) & \text{if } s = s_1 \\ \frac{1}{\hat{\pi}^1(s_1|\omega_j)}(1-\varepsilon)\pi^h(s|\omega_j) & \text{if } s \neq s_1; \end{cases}$$

and if the history at stage $T$ is $hs$ with $s \neq s_1$, the seller offers no further experiment at no cost. This mechanism ensures that $\hat{E}^h$ induces the same distribution of final beliefs as that induced by $e^h$.

It is clear that even in the case of general value functions, the IIR constraints must bind for at least one type at every history, as shown in the proof of Lemma 6. Suppose the IIR constraint at this history $h$ binds for type $i$. Then we have $p_L^h = VI(\theta_i^h; e^h)$ in the original mechanism. After adding another stage to the original mechanism, the seller continues to bind (IIR-hi) by setting

$$\hat{p}_L^h + \left(\varepsilon\phi_i^h(s_1) + 1 - \varepsilon\right)\hat{p}_L^{hs_1} = VI(\theta_i^h; \hat{E}_L^h) = VI(\theta_i^h; e^h) = p_L^h. \tag{13}$$

This $T$-stage mechanism yields a higher payment gap if

$$(\gamma_0^h - \gamma_1^h)\hat{p}_L^{h'} + \left(\gamma_0^h\hat{\pi}^1(s_1|\omega_0) - \gamma_1^h\hat{\pi}^1(s_1|\omega_1)\right)\hat{p}_L^{hs_1} > (\gamma_0^h - \gamma_1^h)p_L^{h'}.$$

Substituting (13), this condition becomes

$$\left(\gamma_0^h\hat{\pi}^1(s_1|\omega_0) - \gamma_1^h\hat{\pi}^1(s_1|\omega_1)\right)\hat{p}_L^{hs_1} > (\gamma_0^h - \gamma_1^h)\left(\varepsilon\phi_i^h(s_1) + 1 - \varepsilon\right)\hat{p}_L^{hs_1},$$

which further simplifies to

$$\gamma_0^h\pi^h(s_1|\omega_0) - \gamma_1^h\pi^h(s_1|\omega_1) > \phi_i^h(s_1)(\gamma_0^h - \gamma_1^h).$$

The last condition holds because $\pi^h(s_1|\omega_0) > \phi_i^h(s_1) > \pi^h(s_1|\omega_1)$. This shows that adding one more stage to a mechanism with $T-1$ stages strictly increases the size of the payment gap and hence relaxes the high type's IC constraint. The seller can strictly increase her revenue by either improving the informativeness of $E_L$ and raising $p_L$, or increasing $p_H$ if $E_L$ is fully informative. ■

## 5. Conclusion

When a data seller can offer a menu of contracts where each contract consists of a sequence of experiments with an associated sequential payment schedule, she can obtain an expected revenue strictly higher than what she could get if each contract were restricted to one experiment at a single price. Our example in the introduction shows that a mechanism involving two-stage experiments can lead to a nontrivial increase in revenue compared to the optimal static mechanism. In this paper, we fully characterize the form of the optimal experiment sequence and payment schedule. We show that if the initial beliefs of the two buyers types are sufficiently far apart, a properly designed dynamic mechanism may allow the seller to fully extract the buyer's surplus despite the presence of private information. When full surplus extraction is not feasible, lengthening the experiment sequence and backloading the payments always leads to a strict revenue gain. Moreover, these insights remain valid in binary state models with general value functions.

It is worth pointing out that the source of the revenue gain does not directly come from a richer space of feasible experiments. Any experiment sequence can be reduced to a one-shot experiment, in the sense that they induce the same distribution of final actions in each state. Our expanded contract space delivers a higher payoff to the seller because selling experiments sequentially allows payments to be contingent on whether the experiment sequence will continue or not. To the extent that different types of buyers attach different probabilities that the experiment sequence will continue, this drives a wedge between their expected costs of the same payment schedule. A mechanism that maximizes this payment gap provides an additional instrument (beyond distorting the quality or quantity of the good being sold) to relax the incentive compatibility constraint, leading to a higher payoff than what is achievable in an optimal static mechanism. Moreover, the optimal extent of the screening distortion becomes weakly smaller when the mechanism induces a positive payment gap.

Taking this perspective further afield, consider an environment in which a buyer's valuation for a durable good is $u(x, \theta)$, where $x$ stands for the degree of durability of the good and $\theta$ denotes buyer type—buyers with higher values of $\theta$ are more patient. A standard model would analyze how the seller can vary $x$ to optimally screen buyers, presumably providing less patient types a less durable product so as to reduce the information rent of more patient types. A moment's reflection would suggest that these different types of buyers may have different preferences over the timing of transfers as well. A less patient buyer, for example, may have stronger preferences for delayed payments compared to more patient types. If this is the case, then offering different payment schedules (payment by installments, or charges for post-sale services, for instance) can potentially serve as an additional screening device to separate different types of buyers. It is therefore theoretically possible that the seller may be able to do strictly better if she can offer a richer set of contracts, in which transfers are not constrained to be paid immediately in a lump sum. Re-examining optimal mechanism design for this and other related screening problems under this light can be a fruitful area of further research.

## Appendix

*Proof of Lemma 1.* (i) We first show that it is without loss of generality to restrict attention to mechanisms $(M_L, M_H)$ such that $C(\theta_{i'}; M_i) \geq C(\theta_i; M_i)$ for $i' \neq i$. Suppose to the contrary that $M_i = (E_i, P_i)$ and $C(\theta_{i'}; M_i) < C(\theta_i; M_i)$. Then replacing $M_i$ with a static contract $(\hat{e}_i, \hat{p}_i)$ and keeping $M_{i'}$ unchanged, where $\hat{e}_i$ satisfies $VI(\theta_i; \hat{e}_i) = VI(\theta_i; E_i)$ and $\hat{p}_i = C(\theta_i; M_i)$, would relax (IC-$i'$) without violating the other constraints.

Next, we show that the fully informative experiment must be part of an optimal mechanism. Suppose to the contrary that both $E_H$ and $E_L$ are imperfect. Without loss of generality, also let $C(\theta_i; M_i) \geq C(\theta_{i'}; M_{i'})$. Then, the seller can modify $E_i$ by replacing all the terminal-stage experiments with fully informative ones and raising the stage-1 price $p_i$ by a positive amount $\varepsilon$, while keeping all other stage prices and stage experiments unchanged. Let $\hat{M}_i = (\hat{E}_i, \hat{P}_i)$ denote the modified mechanism. For $\varepsilon$ small enough, this modification will raise revenue from type $i$ without violating (IC-$i$) or (IR-$i$). The revenue from type $i'$ is unchanged if (IC-$i'$) still holds. If (IC-$i'$) fails under the modification, the revenue from type $i'$ is

$$C(\theta_{i'}; \hat{M}_i) = C(\theta_{i'}; M_i) + \varepsilon > C(\theta_i; M_i) \geq C(\theta_{i'}; M_{i'}).$$

Thus, this modification would raise the seller's revenue.

The above argument shows that at least one type must purchase the fully informative experiment. Suppose that $E_L$ is fully informative but $E_H$ is imperfect. Then, the price that can be charged to the low type is at most $VI(\theta_L)$. Therefore, (IC-$H$) implies that the expected total payment for the high type must be smaller than $VI(\theta_L)$. If this is the case, however, the seller can increase her revenue by assigning the fully informative experiment to the high type as well, resulting in a contradiction.

(ii) Clearly, at least one of (IR-$L$) and (IR-$H$) must bind. Suppose that (IR-$H$) binds but (IR-$L$) does not. Since the high type purchases the fully informative experiment, the seller charges $VI(\theta_H)$ to the high type. Then, (IC-$L$) is given by

$$VI(\theta_L; E_L) - C(\theta_L; M_L) \geq VI(\theta_L) - VI(\theta_H).$$

The right-hand side of this inequality is strictly negative by Assumption 1. Thus, the seller can raise revenue without violating any constraint by raising, say, the stage-1 price $p_L$.

(iii) If neither (IC-$H$) nor (IR-$H$) binds, then the seller could obviously raise revenue by charging more to the high type. ■

*Proof of Lemma 2.* Define $\underline{\mathscr{H}} := \{h \in \mathscr{H} : C(\theta_H^h; M_L^h) > VI(\theta_H^h; E_L^h)\}$ as the set of histories at which (IIR-$hH$) is violated. Towards a contradiction, we consider a mechanism $M_L = (E_L, P_L)$ in which $\underline{\mathscr{H}}$ is nonempty, and construct a modified mechanism that is weakly better. There are two cases.

*Case 1:* $E_L$ is not fully informative. Consider a modified contract $\hat{M}_L = (\hat{E}_L, \hat{P}_L)$ in which $\hat{E}_L$ is fully informative with probability $\varepsilon$ and the same as $E_L$ with probability $1-\varepsilon$, where $\varepsilon$ is some arbitrarily small positive number. This modification increases the value of information for the low type by

$$VI(\theta_L; \hat{E}_L) - VI(\theta_L; E_L) = \varepsilon(VI(\theta_L) - VI(\theta_L; E_L)) > 0.$$

We then set $\hat{p}_L = p_L + VI(\theta_L; \hat{E}_L) - VI(\theta_L; E_L)$ while keeping everything else unchanged. If (IC-$H$) is slack, this modification strictly raises the low type's payment while still satisfying (IC-$H$). Now suppose (IC-$H$) binds, in which case the change in the high type's information rent is

$$\varepsilon(VI(\theta_H) - VI(\theta_H; E_L)) - \varepsilon(VI(\theta_L) - VI(\theta_L; E_L)). \tag{14}$$

If (14) is weakly negative, then (IC-$H$) still holds. If (14) is strictly positive, we can increase $\hat{p}_L$ by this amount and reduce $\hat{p}_L^h$ at some $h \in \underline{\mathscr{H}}$ in a way to keep $C(\theta_L; \hat{M}_L)$ unchanged. This modification raises the low type's expected payment without raising the high type's information rent.

*Case 2:* $E_L$ is fully informative. If (IR-$H$) is slack, we set $\hat{p}_L = p_L + \varepsilon$ and reduce $\hat{p}_L^h$ at some $h \in \underline{\mathscr{H}}$ in a way to keep $C(\theta_L; \hat{M}_L)$ unchanged. This modification raises $C(\theta_H; E_L)$ and strictly reduces the high type's information rent, which is a strict improvement. If (IR-$H$) binds, the seller extracts full surplus and cannot reduce the high type's information rent any further. In this case, we consider an alternative mechanism that extracts full surplus while satisfying (IIR-$hH$) at all $h \in \mathscr{H}$. Specifically, construct a modified contract $\hat{M}_L = (E_L, \hat{P}_L)$ by: (i) reducing $\hat{p}_L^h$ to satisfy $VI(\theta_H^h) = C(\theta_H^h; \hat{M}_L)$ at all $h \in \underline{\mathscr{H}}$, while keeping $\hat{p}_L^h = p_L^h$ at all $h \in \mathscr{H} \backslash \underline{\mathscr{H}}$; and (ii) raising $\hat{p}_L$ to satisfy $C(\theta_L; \hat{M}_L) = C(\theta_L; M_L)$. Note that the expected continuation net payoff for the high type at any $h \in \underline{\mathscr{H}}$ is 0 in the original mechanism, while it changes to $VI(\theta_H^h) - C(\theta_H^h; \hat{M}_L^h)$ after the modification, which is also 0 by construction. Therefore, this modification does not affect the high type's information rent, meaning that the modified mechanism continues to achieve full surplus extraction while satisfying (IIR-$hH$) at all $h \in \mathscr{H}$. ∎

**Claim 1.** *In an optimal mechanism, $p_L^{hs} = 0$ for any history $h \in \mathscr{H}_1 \cup \mathscr{H}$ followed by an up signal $s \in S^h$.*

*Proof.* Suppose to the contrary that $p_L^{hs} > 0$. The seller can reduce $p_L^{hs}$ by $\varepsilon > 0$ and raise $p_L^h$ by $(\gamma_0^{hs}/\gamma_0^h)\varepsilon$. From equation (5), this operation raises the payment gap if

$$(\gamma_0^h - \gamma_1^h)\frac{\gamma_0^{hs}}{\gamma_0^h}\varepsilon \geq (\gamma_0^{hs} - \gamma_1^{hs})\varepsilon,$$

which is true because the likelihood ratios satisfy $\gamma_1^{hs}/\gamma_0^{hs} \geq \gamma_1^h/\gamma_0^h$ for any up signal $s$. Observe also that this operation changes the expected payment for type $i$ at history $h$ by

$$\frac{\gamma_0^{hs}}{\gamma_0^h}\varepsilon - \left[\theta_i^{hs}\frac{\gamma_1^{hs}}{\gamma_1^h} + (1-\theta_i^{hs})\frac{\gamma_0^{hs}}{\gamma_0^h}\right]\varepsilon \leq 0.$$

Thus, (IIR-$hi$) will be satisfied for both types. ∎

*Proof of Lemma 4.* By Claim 1, the seller does not charge a positive price at any history ending with an up signal. Eliminating such a history will not affect the seller's revenue or the payment gap. Specifically, suppose that history $hu$ is not a terminal history and $u \in S^h$ is an up signal. Let $\{\theta_i^{hus'}\}_{s' \in S^{hu}}$ be the set of possible beliefs at the next stage following $hu$ in the original experiment sequence $E_L$. We may modify the experiment sequence $E_L$ to $\hat{E}_L$ in such a way that at history $h$, we eliminate the up history $hu$ by setting $\hat{\pi}^h(u|\omega_j) = 0$ for $j = 0, 1$ (and keeping $\hat{\pi}^h(s|\omega_j)$ unchanged for each $s \in S^h$, $s \neq u$). We allocate this probability by adding signals $s'$ in the set $S^{hu}$ to the set $\hat{S}^h$ and setting $\hat{\pi}^h(s'|\omega_j) = \pi^h(u|\omega_j)\pi^{hu}(s'|\omega_j)$ for each $s' \in S^{hu}$. At any history $hs'$ introduced by this modification, the seller offers an uninformative experiment at stage $t$ at zero price. This modification would eliminate history $hu$ without altering the seller's revenue or the payment gap. Figure 3 graphically illustrates how we can construct such an equivalent mechanism. This argument shows that it is without loss of generality to assume that any history $hu$ is a terminal history if $u$ is an up signal. Furthermore, if the action chosen at such terminal history $hu$ is $a_0$, this would imply that $a_0$ occurs with positive probability in state $\omega_1$, in contradiction to Lemma 3. ∎

**Claim 2.** *At any history $hs$ where $s$ is a down signal, $\theta_H^{hs} < \sigma$.*

*Proof.* Suppose $M_L = (E_L, P_L)$ is an optimal mechanism that would induce $\theta_H^{hs} \geq \sigma$ at history $hs$ where $s$ is a down signal. This implies $\theta_L^{hs} > \sigma$. We can construct a modified contract $\hat{M}_L = (\hat{E}_L, \hat{P}_L)$ that strictly raises the seller's revenue. The modified contract splits $\theta_L^{hs}$ into $\theta_L^{hs^-}$ and $\theta_L^{hs^+}$, with $\theta_L^{hs^-} < \theta_L^{hs} < \theta_L^{hs^+}$. Specifically, we choose $\hat{\gamma}_j^h$ such that

$$\hat{\gamma}_1^{hs^-} = \gamma_1^{hs} - \varepsilon, \quad \hat{\gamma}_0^{hs^-} = \gamma_0^{hs}, \quad \hat{\gamma}_1^{hs^+} = \varepsilon, \quad \hat{\gamma}_0^{hs^+} = 0.$$

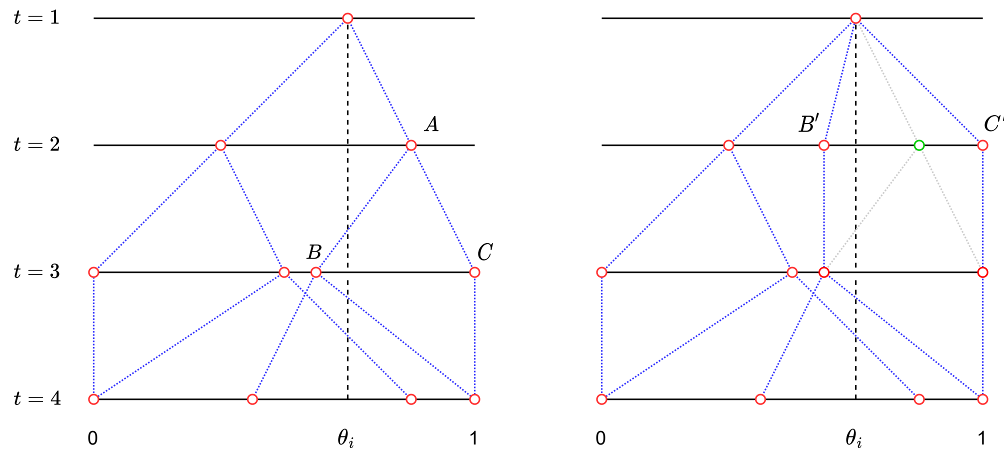


***Figure 3.*** *An equivalent mechanism. In the left panel, the node A at $t=2$ is a history followed by an up signal but is not a terminal history. We can eliminate A and bring beliefs B and C to $B'$ and $C'$, respectively, to $t=2$ to construct an equivalent mechanism, as depicted in the right panel.*

For $\varepsilon$ positive and sufficiently small, we have $\theta_i^{hs^-} \in (\sigma, \theta_i^{hs})$ and $\theta_i^{hs^+} = 1$ for both $i = L, H$. Following $hs^-$, we consider the same signal space, i.e., $\hat{S}^{hs^-} = S^{hs} = \{s_1, \ldots, s_N, u\}$, where $u$ stands for an up signal. We choose

$$\hat{\gamma}_0^{hs^-s'} = \gamma_0^{hss'} \quad \text{for all } s' \in S^{hs},$$
$$\hat{\gamma}_1^{hs^-s'} = \gamma_1^{hss'} \quad \text{for } s' = s_1, \ldots, s_N, \qquad \hat{\gamma}_1^{hs^-u} = \gamma_1^{hsu} - \varepsilon.$$

If the state is $\omega_1$, this modification will move some probability mass from the history $hsu$ to the new history $hs^+$. By Lemma 4, both histories are terminal histories that result in the same action $a_1$, and therefore the distribution of final actions in each state remains unaffected. Hence, at any history preceding (and including) $hs$, the gross payoff from the modified experiment sequence is the same as that from the original sequence for both types of buyer.

If the seller charges the same price at history $hs^-$ in the modified contract, i.e., $\hat{p}_L^{hs^-} = p_L^{hs}$, then the payment gap will widen, because

$$(\hat{\gamma}_0^{hs^-} - \hat{\gamma}_1^{hs^-})\hat{p}_L^{hs^-} = (\gamma_0^{hs} - \gamma_1^{hs} + \varepsilon)p_L^h > (\gamma_0^{hs} - \gamma_1^{hs})p_L^{hs}.$$

Since a larger payment gap leads to a strict improvement in revenue, this would imply that the original mechanism cannot be optimal, unless setting $\hat{p}_L^{hs^-} = p_L^{hs}$ is infeasible. Therefore, it suffices to show that $VI(\theta_i^{hs^-}; \hat{E}_L^{hs^-}) \geq VI(\theta_i^{hs}; E_L^{hs})$, so that the seller can charge the same price at $hs^-$ without violating the IIR constraints at $hs^-$.

By construction, the modified contract splits $\theta_i^{hs}$ into $\theta_i^{hs^-}$ and $\theta_i^{hs^+} = 1$, without changing the gross payoff from the experiment. Therefore,

$$\mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^{hs}, E_L^{hs}\right] = \left(1 - \hat{\phi}_i^h(1)\right)\mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^{hs^-}, \hat{E}_L^{hs^-}\right] + \hat{\phi}_i^h(1)v(1),$$

where $\hat{\phi}_i^h(1)$ is type $i$'s probability assessment at history $h$ that the belief will be split to $\theta_i^{hs^+} = 1$. From this, we obtain

$$\begin{aligned}
&\left(\mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^{hs^-}, \hat{E}_L^{hs^-}\right] - v(\theta_i^{hs^-})\right) - \left(E\left[v(\tilde{\theta}_i) \mid \theta_i^{hs}, E_L^{hs}\right] - v(\theta_i^{hs})\right) \\
&= \hat{\phi}_i^h(1)\left(\mathbb{E}\left[v(\tilde{\theta}_i) \mid \theta_i^{hs^-}, \hat{E}_L^{hs^-}\right] - v(\theta_i^{hs^-})\right) + v(\theta_i^{hs}) - \left(1 - \hat{\phi}_i^h(1)\right)v(\theta_i^{hs^-}) - \hat{\phi}_i^h(1)v(1) \\
&\geq v(\theta_i^{hs}) - \left(1 - \hat{\phi}_i^h(1)\right)v(\theta_i^{hs^-}) - \hat{\phi}_i^h(1)v(1).
\end{aligned}$$

Recall that we assume $\sigma \leq \theta_i^{hs^-} < \theta_i^{hs} < 1$. Since the beliefs $\theta_i^{hs^-}$ and 1 are a mean-preserving spread of $\theta_i^{hs}$, and $v(\cdot)$ is affine on $[\sigma, 1]$, the last expression is equal to 0. This shows that $VI(\theta_i^{hs^-}; \hat{E}_L^{hs^-}) \geq VI(\theta_i^{hs}; E_L^{hs})$ for $i = L, H$. Thus, the modified mechanism is feasible and strictly improves on the original mechanism, which is a contradiction. ■

*Proof of Lemma 5.* By Claim 2, $\theta_H^{hs} < \sigma$ whenever $s$ is a down signal. Toward a contradiction, suppose $hs$ is a terminal history with $\theta_H^{hs} > 0$. Then the high type would choose action $a_0$ at $hs$. Note that $\theta_H^{hs} > 0$ implies that history $hs$ is reached with positive probability in state $\omega_1$. This contradicts $\Gamma_{1a_0} = 0$ established in Lemma 3. ■

*Proof of Lemma 6.* (i) We first show that at every history, at least one of the IIR constraints must bind. Suppose to the contrary that there is some $h$ following a predecessor history $h'$, such that the (IIR-$hi$) is slack for both $i = L, H$. By Lemma 4, it is without loss of generality to assume that $\theta_i^{h'} > \theta_i^h$ and $p_L^{h'} > 0$, for otherwise (IR-$L$) would be slack, leading to a contradiction. We can then lower $p_L^{h'}$ by $(\gamma_0^h/\gamma_0^{h'})\varepsilon$ and raise $p_L^h$ by $\varepsilon$. The change in the expected payment for type $i$ at $h'$ is

$$-\frac{\gamma_0^h}{\gamma_0^{h'}}\varepsilon + \left[\theta_i^{h'}\frac{\gamma_1^h}{\gamma_1^{h'}} + (1 - \theta_i^{h'})\frac{\gamma_0^h}{\gamma_0^{h'}}\right]\varepsilon = \theta_i^{h'}\left(\frac{\gamma_1^h}{\gamma_1^{h'}} - \frac{\gamma_0^h}{\gamma_0^{h'}}\right)\varepsilon < 0,$$

because $\theta_i^{h'} > \theta_i^h$. This ensures that (IIR-$h'i$) still holds for both types. On the other hand, the change in the payment gap is proportional to

$$-(\gamma_0^{h'} - \gamma_1^{h'})\frac{\gamma_0^h}{\gamma_0^{h'}} + (\gamma_0^h - \gamma_1^h) > 0.$$

This operation is thus a strict improvement over the original mechanism.

Given that at least one of the IIR constraints must bind at each history, to establish part (i), it suffices to rule out the possibility that (IIR-$h'L$) binds and (IIR-$h'H$) is slack at some history $h'$. Suppose such a $h'$ exists. There are two possibilities.

*Case 1:* $\theta_L^{h'} > \sigma$. In this case, we can construct a modified contract $\hat{M}_L = (\hat{E}_L, \hat{P}_L)$ in the same manner as is described in the proof of Claim 2. The modified contract would split the belief $\theta_L^{h'}$ into $\theta_L^{h-}$ and $\theta_L^{h+} = 1$, leading to a larger payment gap than the original contract $M_L = (E_L, P_L)$ without affecting the gross payoff from the experiment. Moreover, the proof of Claim 2 shows that

$$VI(\theta_i^{h^-}; \hat{E}_L^{h^-}) - VI(\theta_i^{h'}; E_L^{h'}) \geq v(\theta_i^{h'}) - \left(1 - \hat{\phi}_i^{h'}(1)\right) v(\theta_i^{h^-}) - \hat{\phi}_i^{h'}(1) v(1).$$

For $i = L$, we can choose $\theta_L^{h^-}$ sufficiently close to $\theta_L^{h'}$ such that $\theta_L^{h'} > \theta_L^{h^-} > \sigma$. Since $v(\cdot)$ is affine on $[\sigma, 1]$, the right-hand side of the inequality is equal to 0. Therefore, $VI(\theta_L^{h^-}; \hat{E}_L^{h^-}) \geq VI(\theta_L^{h'}; E_L^{h'})$, and we can set the same price, $\hat{p}_L^{h^-} = p_L^{h'}$, without violating (IIR-$h^-L$). For $i = H$, $VI(\theta_H^{h^-}; \hat{E}_L^{h^-})$ will be arbitrarily close to $VI(\theta_H^{h'}; E_L^{h'})$ for $\theta_H^{h^-}$ close to $\theta_H^{h'}$. Therefore, if (IIR-$h'H$) holds with slack, (IIR-$h^-H$) will hold by maintaining the same price $\hat{p}_L^{h^-} = p_L^{h'}$. This shows that the modified mechanism is feasible and increases the seller's revenue, which is a contradiction.

*Case 2:* $\theta_L^{h'} \leq \sigma$. This case implies $v(\theta_H^{h'}) = (1 - \theta_H^{h'})\beta$ and $v(\theta_L^{h'}) = (1 - \theta_L^{h'})\beta$. Let $\Gamma_{0a_0}^{h'}$ be the probability of choosing $a_0$ in state $\omega_0$ at history $h'$. It follows from equation (4) that

$$\Delta VI(E_L) = -(\theta_L^{h'} - \theta_H^{h'})(1 - \Gamma_{0a_0}^{h'})\beta \leq 0.$$

On the other hand, Claim 1 implies that the payment gap must satisfy $\Delta C(M_L) \geq 0$. These two conditions contradict the premise that (IIR-$h'L$) binds while (IIR-$h'H$) is slack.

(ii) Consider a history $h \in \mathscr{H}$ followed by signal $s \in S^h$. Part (i) of this lemma shows that $C(\theta_H^{hs}; M_L^{hs}) = VI(\theta_H^{hs}; E_L^{hs})$ for all $s$. From this, we obtain

$$\begin{aligned}
\sum_{s \in S^h} \phi_H^h(s) C(\theta_H^{hs}; M_L^{hs}) &= \sum_{s \in S^h} \phi_H^h(s) VI(\theta_H^{hs}; E_L^{hs}) \\
&= \sum_{s \in S^h} \phi_H^h(s) \left( \mathbb{E}\left[ v(\tilde{\theta}_H) \mid \theta_H^{hs}; E_L^{hs} \right] - v(\theta_H^{hs}) \right), \\
&= VI(\theta_H^h; E_L^h) + v(\theta_H^h) - \sum_{s \in S^h} \phi_H^h(s) v(\theta_H^{hs}),
\end{aligned}$$

where the second equality follows from (1), and the third follows from the recursion (2). The left-hand side of the above equation plus the stage price $p_L^h$ is equal to $C(\theta_H^h; M_L^h)$.

Using the fact that $C(\theta_H^h; M_L^h) = VI(\theta_H^h; E_L^h)$, we obtain

$$p_L^h = \sum_{s \in S^h} \phi_H^h(s) v(\theta_H^{hs}) - v(\theta_H^h) = VI(\theta_H^h; e_L^h). \qquad \blacksquare$$

**Claim 3.** *For any terminal history $h \in \mathscr{H}_t$ ending with an up signal at stage $t \geq 3$, $\theta_i^h = 1$ for $i = L, H$.*

*Proof.* Consider some history $h \in \mathscr{H}$ with $S^h = \{s_1, \ldots, s_N, u\}$, where we generically denote an up signal by $u$. Also let $S^{hs_1} = \{s_1', \ldots, s_{N'}', u'\}$. The (unconditional) payment gap over the two stages in this particular sequence of events (i.e., $h$ followed by $s_1$) is

$$(\theta_L - \theta_H)(\gamma_0^h - \gamma_1^h)\left[(1-\theta_H^h)\sum_{n=1}^{N} \frac{\gamma_0^{hs_n}}{\gamma_0^h}\beta + \theta_H^h \frac{\gamma_1^{hu}}{\gamma_1^h}\alpha - (1-\theta_H^h)\beta\right]$$
$$+ (\theta_L - \theta_H)(\gamma_0^{hs_1} - \gamma_1^{hs_1})\left[(1-\theta_H^{hs_1})\sum_{n=1}^{N'} \frac{\gamma_0^{hs_1 s_n'}}{\gamma_0^{hs_1}}\beta + \theta_H^{hs_1} \frac{\gamma_1^{hs_1 u'}}{\gamma_1^{hs_1}}\alpha - (1-\theta_H^{hs_1})\beta\right],$$

where we use $p_L^{h'} = VI(\theta_H^{h'}, e_L^{h'})$ and $\theta_H^{h'} < \sigma$ for histories $h' = h, hs_1$, established in Lemma 6 and Claim 2. By Bayes' rule, this is equal to

$$\frac{(\theta_L - \theta_H)}{\theta_H \gamma_1^h + (1-\theta_H)\gamma_0^h}(\gamma_0^h - \gamma_1^h)\left[(1-\theta_H)\sum_{n=1}^{N} \gamma_0^{hs_n}\beta + \theta_H \gamma_1^{hu}\alpha - (1-\theta_H)\gamma_0^h\beta\right]$$
$$+ \frac{(\theta_L - \theta_H)}{\theta_H \gamma_1^{hs_1} + (1-\theta_H)\gamma_0^{hs_1}}(\gamma_0^{hs_1} - \gamma_1^{hs_1})\left[(1-\theta_H)\sum_{n=1}^{N'} \gamma_0^{hs_1 s_n'}\beta + \theta_H \gamma_1^{hs_1 u'}\alpha - (1-\theta_H^{hs_1})\gamma_0^{hs_1}\beta\right]$$
$$= \left[1 - \frac{\mathbb{P}[h \mid \theta_L, E_L]}{\mathbb{P}[h \mid \theta_H, E_L]}\right]\left[\theta_H \gamma_1^{hu}\alpha - (1-\theta_H)\gamma_0^{hu}\beta\right]$$
$$+ \left[1 - \frac{\mathbb{P}[hs_1 \mid \theta_L, E_L]}{\mathbb{P}[hs_1 \mid \theta_H, E_L]}\right]\left[\theta_H \gamma_1^{hs_1 u'}\alpha - (1-\theta_H)\gamma_0^{hs_1 u'}\beta\right],$$

where $\mathbb{P}[h' | \theta, E] = \theta \gamma_1^{h'} + (1-\theta)\gamma_0^{h'}$ is the ex-ante probability under initial belief $\theta$ that experiment sequence $E$ will yield history $h'$, for $h' = h, hs_1$.

Suppose for the sake of contradiction that $\gamma_0^{hs_1 u'} > 0$ at stage $t \geq 3$. We consider a modified contract which:

- decrease $\gamma_0^{hs_1 u'}$ by $\varepsilon$ and decrease $\gamma_1^{hs_1 u'}$ by $(\gamma_1^{hs_1}/\gamma_0^{hs_1})\varepsilon$ (this will increase the likelihood ratio of history $hs_1 u'$ and raise $p_L^{hs_1}$);
- increase $\gamma_0^{hu}$ by $\varepsilon$ and increase $\gamma_1^{hu}$ by $(\gamma_1^{hs_1}/\gamma_0^{hs_1})\varepsilon$ (this will decrease the likelihood ratio of history $hu$ and lower $p_L^h$);

- decrease $\gamma_0^{hs_1}$ by $\varepsilon$ and decrease $\gamma_1^{hs_1}$ by $(\gamma_1^{hs_1}/\gamma_0^{hs_1})\varepsilon$ (this will keep the likelihood ratio of $hs_1$ unchanged and hence $\mathbb{P}[hs_1|\theta_L, E_L]/\mathbb{P}[hs_1|\theta_H, E_L]$ will remain unchanged).

The marginal change in the payment gap from this modification is

$$\left[\frac{\mathbb{P}[hs_1 \mid \theta_L, E_L]}{\mathbb{P}[hs_1 \mid \theta_H, E_L]} - \frac{\mathbb{P}[h \mid \theta_L, E_L]}{\mathbb{P}[h \mid \theta_H, E_L]}\right]\left[\theta_H \alpha \frac{\gamma_1^{hs_1}}{\gamma_0^{hs_1}} \varepsilon - (1-\theta_H)\beta \varepsilon\right].$$

Because $s_1$ is a down signal, we have $\gamma_1^{hs_1}/\gamma_0^{hs_1} < \gamma_1^h/\gamma_0^h$, which implies that the first term in the above is negative. The second term in the above has the same sign as $\theta_H^{hs_1} - \sigma$, which is also negative from Claim 2. Thus the modified contract strictly increases the payment gap. This shows that we must have $\gamma_0^{h'} = 0$ for any history $h'$ ending with an up signal at stage $t \geq 3$, which in turn implies $\theta_i^{h'} = 1$ for both types $i$ at such history. ■

*Proof of Proposition 1.* Claim 3 shows that $\theta_i^h = 1$ if $h \in \mathscr{H}_{T+1}$ ends with an up signal at the final stage, while Lemma 5 shows that $\theta_i^h = 0$ if $h \in \mathscr{H}_{T+1}$ ends with a down signal. This implies that any stage-$T$ experiment is fully informative and binary.

We next show that it is without loss of generality to assume that, in an optimal mechanism, any experiment at stage $t = 1, \ldots, T-1$ is binary. Suppose the original mechanism is $M_L = (E_L, P_L)$, and there is some history $h \in \mathscr{H}_t$ with $e_L^h = (S^h, \pi^h)$. Let $S^h = \{s_1, \ldots, s_N, u\}$, where $s_1, \ldots, s_N$ are multiple down signals and $u$ is an up signal. Consider $N$ alternative mechanisms, $\hat{M}_n = (\hat{E}_n, \hat{P}_n)$ for $n = 1, \ldots, N$, where $\hat{E}_n$ is identical to $E_L$ except that at history $h$, the experiment $e_L^h$ is replaced by a binary experiment $\hat{e}_n = (\{d_n, u_n\}, \hat{\pi}_n)$ where, for $j = 0, 1$,

$$\hat{\pi}_n(d_n|\omega_j) = \frac{1}{\lambda_n}\pi^h(s_n|\omega_j),$$

with

$$\lambda_n = \frac{\pi^h(u|\omega_1)\pi^h(s_n|\omega_0) - \pi^h(s_n|\omega_1)\pi^h(u|\omega_0)}{\pi^h(u|\omega_1) - \pi^h(u|\omega_0)}.$$

Note that $\lambda_n$ is positive and $\sum_{n=1}^N \lambda_n = 1$. Moreover, by construction, the likelihood ratio of $d_n$ under $\hat{e}_n$ is the same as the likelihood ratio of $s_n$ under $e_L^h$. These two signals would induce the same posterior beliefs, i.e., $\hat{\theta}_i^{hd_n} = \theta_i^{hs_n}$, in the respective mechanisms. Moreover, it can be verified that, for any $n$, the likelihood ratio $u_n$ under $\hat{e}_n$ is the same as the likelihood ratio of $u$ under $e_L^h$; thus, the respective mechanisms also yield the same posterior belief upon up signals, i.e., $\hat{\theta}_H^{hu_n} = \theta_H^{hu}$.

The payment schedule $\hat{P}_n$ is the same as $P_L$ except at history $h$, where we set

$$\begin{aligned}\hat{p}_n^h = VI(\theta_H^h; \hat{e}_n) &= \sum_{s\in\{d_n,u_n\}} \hat{\phi}_H^h(s)v(\hat{\theta}_H^{hs}) - v(\theta_H^h)\\ &= \frac{1}{\lambda_n}\phi_H^h(s_n)v(\theta_H^{hs_n}) + \left[1-\frac{1}{\lambda_n}\phi_H^h(s_n)\right]v(\theta_H^{hu}) - v(\theta_H^h).\end{aligned}$$

Clearly, we have $\sum_{n=1}^N \lambda_n \hat{p}_n^h = p_L^h$. Let $\mathscr{H}^{hs_n}$ be the set of histories that can be reached with positive probability after history $hs_n$. Then the payment gap induced by mechanism $\hat{M}_n$ is

$$\Delta C(\hat{M}_n) = (\theta_L - \theta_H)\left[\sum_{h'\in\mathscr{H}_\tau\setminus\{h\},\tau\le t} (\gamma_1^{h'} - \gamma_0^{h'})p_L^{h'} + (\gamma_1^h - \gamma_0^h)\hat{p}_n^h + \sum_{h''\in\mathscr{H}^{hs_n}} \frac{1}{\lambda_n}(\gamma_1^{h''} - \gamma_0^{h''})p_L^{h''}\right].$$

We therefore obtain

$$\Delta C(M_L) = \sum_{n=1}^N \lambda_n \Delta C(\hat{M}_n).$$

Let $n^* \in \arg\max_n \Delta C(\hat{M}_n)$. Then we have $\Delta C(\hat{M}_{n^*}) \ge \Delta C(M_L)$. Also, $\hat{M}_{n^*}$ and the original mechanism $M_L$ produce the same distribution of final actions in each state. Thus, replacing $M_L$ by $\hat{M}_{n^*}$ would enlarge the payment gap without altering the informativeness of the experiment sequence.

Furthermore, if the original mechanism satisfies the IIR constraints, so does $\hat{M}_{n^*}$. To see this, it suffices to consider the IIR constraints at history $h \in \mathscr{H}$. By construction, $\hat{p}_{n^*}^h$ is chosen to bind (IIR-$hH$). For (IIR-$hL$), observe that the continuation experiment sequences at history $h$ under $E_L$ and $\hat{E}_{n^*}$ are both fully informative. Thus,

$$VI(\theta_i^h; E_L) = VI(\theta_i^h; \hat{E}_{n^*}). \tag{15}$$

Applying (15) to $i = H$, and using the fact that (IIR-$hH$) binds for both the original and modified mechanisms, we obtain

$$C(\theta_H^h; M_L^h) = C(\theta_H^h; \hat{M}_{n*}^h).$$

Together with the fact that $\Delta C(\hat{M}_{n^*}) \ge \Delta C(M_L)$, this implies

$$C(\theta_L^h; M_L^h) \ge C(\theta_L^h; \hat{M}_{n^*}^h).$$

Since (15) also applies to $i = L$, the above inequality implies that if (IIR-$hL$) holds in $M_L^h$, it must hold in $\hat{M}_{n^*}^h$. This shows that $\hat{M}_{n^*}$ is feasible, and therefore the original mechanism is not optimal. This completes the proof that it is without loss of generality to consider binary mechanisms.

Part (i) of the proposition then follows because, in a binary experiment at stage $\tau \geq 2$, $\pi^t_{0d} < 1$ would imply that an up history $h'$ at stage $\tau+1 \geq 3$ occurs with positive probability in state $\omega_0$. This contradicts $\theta^{h'}_i = 1$ required by Claim 3. Part (ii) follows because $\pi^T_{1u} < 1$ would contradict $\Gamma_{1a_0} = 0$ required by Lemma 3. ■

**Claim 4.** *Let*

$$G^2_T(\theta^2_H) := \frac{\alpha}{\theta^2_H}\sum_{t=3}^{T}\frac{\theta^2_H - \psi^t(\theta^2_H))(\psi^t(\theta^2_H) - \psi^{t+1}(\theta^2_H))}{(1-\psi^t(\theta^2_H))(1-\psi^{t+1}(\theta^2_H))},$$

*where $\psi^t(\theta^2_H) = 1-(1-\theta^2_H)^{\frac{T-t+1}{T-1}}$. The functions $G_T$ (defined in equation (7)) and $G^2_T$ satisfy the following properties:*

*(i) For any $\theta^2_H < \sigma$, $G_T(\theta^2_H;\theta_H) < \beta$ and $G^2_T(\theta^2_H) < \beta$.*

*(ii) $\frac{\partial G_T(\theta^2_H;\theta_H)}{\partial \theta_H} < \frac{1}{1-\theta_H}\alpha$ and $\frac{\partial G^2_T(\theta^2_H)}{\partial \theta^2_H} < \frac{1}{1-\theta^2_H}\alpha$.*

*Proof.* (i) Using the formula for $\psi^t(\cdot)$, we obtain

$$\frac{\psi^\tau(\theta^2_H) - \psi^{\tau+1}(\theta^2_H)}{1-\psi^{\tau+1}(\theta^2_H)} = 1-(1-\theta^2_H)^{\frac{1}{T-1}}.$$

Thus,

$$\begin{aligned} G_T(\theta^2_H;\theta_H) &= \frac{\alpha}{\theta_H}\Big[1-(1-\theta^2_H)^{\frac{1}{T-1}}\Big]\sum_{\tau=2}^{T}\frac{\theta_H - \psi^\tau(\theta^2_H)}{1-\psi^\tau(\theta^2_H)} \\ &< \alpha\Big[1-(1-\theta^2_H)^{\frac{1}{T-1}}\Big]\sum_{\tau=2}^{T}\frac{1}{1-\psi^\tau(\theta^2_H)} \\ &= \alpha\frac{\theta^2_H}{1-\theta^2_H}, \end{aligned}$$

which is less than $\beta$ for $\theta^2_H < \sigma$. Similarly,

$$G^2_T(\theta^2_H) = \frac{\alpha}{\theta^2_H}\Big[1-(1-\theta^2_H)^{\frac{1}{T-1}}\Big]\sum_{t=3}^{T}\frac{\theta^2_H - \psi^t(\theta^2_H)}{1-\psi^t(\theta^2_H)} < \alpha\frac{\theta^3_H}{1-\theta^3_H} < \beta.$$

(ii) The derivative of $G_T$ is

$$\begin{aligned}\frac{\partial G_T(\theta_H^2;\theta_H)}{\partial \theta_H} &= -\frac{G_T(\theta_H^2;\theta_H)}{\theta_H} + \frac{\alpha}{\theta_H}\sum_{t=2}^{T}\frac{\psi^t(\theta_H^2)-\psi^{t+1}(\theta_H^2)}{(1-\psi^t(\theta_H^2))(1-\psi^{t+1}(\theta_H^2))}\\ &< \frac{\alpha}{\theta_H}\sum_{t=2}^{T}\frac{\psi^t(\theta_H^2)-\psi^{t+1}(\theta_H^2)}{(1-\psi^t(\theta_H^2))(1-\psi^{t+1}(\theta_H^2))}\\ &= \frac{\alpha}{\theta_H}\frac{\theta_H^2}{1-\theta_H^2},\end{aligned}$$

which is less than $\alpha/(1-\theta_H)$ because $\theta_H^2 < \theta_H$. Similarly, by the envelope theorem,

$$\frac{\partial G_T^2(\theta_H^2)}{\partial \theta_H^2} = -\frac{G_T^2(\theta_H^2)}{\theta_H^2} + \frac{\alpha}{\theta_H^2}\sum_{t=3}^{T}\frac{\psi^t(\theta_H^2)-\psi^{t+1}(\theta_H^2)}{(1-\psi^t(\theta_H^2))(1-\psi^{t+1}(\theta_H^2))} < \frac{\alpha}{\theta_H^2}\frac{\theta_H^3}{1-\theta_H^3} < \frac{\alpha}{1-\theta_H^2}. \quad \blacksquare$$

*Proof of Lemma 7.* We first establish the existence and uniqueness of $\theta_T^*$ in $(\underline{\theta},\sigma)$. Define $J_T^2(\theta_H^2;\theta_H)$ to be equal to the left-hand side of (9). If $\rho(\theta_H^2;\theta_H) \le \sigma$, then $VI(\rho(\theta_H^2;\theta_H)) > VI(\theta_H^2)$, and (IIR-2$L$) is slack. We thus only need to consider the case where $\rho(\theta_H^2;\theta_H) > \sigma$. This implies that $VI(\rho(\theta_H^2;\theta_H)) = (1-\rho(\theta_H^2;\theta_H))\beta$. Thus, we can write

$$J_T^2(\theta_H^2;\theta_H) = (1-\theta_H^2)\beta - \theta_H^2\alpha - (\rho(\theta_H^2;\theta_H)-\theta_H^2)(\beta - G_T^2(\theta_H^2)).$$

Since $VI(\underline{\theta}) = VI(\theta_L) < VI(\rho(\underline{\theta};\theta_H))$, we have $J_T^2(\underline{\theta};\theta_H) > 0$. Note also that $J_T^2(\sigma;\theta_H) = -(\rho(\sigma;\theta_H)-\sigma)(\beta - G_T^2(\sigma)) < 0$ by part (i) of Claim 4. This shows that there exists $\theta_T^* \in (\underline{\theta},\sigma)$ such that $J_T^2(\theta_T^*;\theta_H) = 0$. Moreover,

$$\frac{\partial J_T^2}{\partial \theta_H^2} = -\alpha - G_T^2(\theta_H^2) + (\rho(\theta_H^2;\theta_H)-\theta_H^2)\frac{\partial G_T^2}{\partial \theta_H^2} - \frac{\partial \rho}{\partial \theta_H^2}(\beta - G_T^2(\theta_H^2)),$$

where $\partial G_T^2/\partial\theta_H^2 < \alpha/(1-\theta_H^2)$ by part (ii) of Claim 4, and $\partial\rho/\partial\theta_H^2 > 0$. Thus, $J_T^2(\cdot;\theta_H)$ is strictly decreasing. It follows that $\theta_T^*$ is unique, and (IIR-2$L$) holds for all $\theta_H^2 \in [\underline{\theta},\theta_T^*]$.

The first-order condition for maximizing $G_T(\theta_H^2;\theta_H)$ is

$$(\theta_H - \theta_H^2)(1-\theta_H^2) - (\theta_H^2 - \psi^3(\theta_H^2))(1-\theta_H) = 0.$$

It can readily be verified that $\theta_T^{**} = 1-(1-\theta_H)^{\frac{T-1}{T}}$ solves the first-order condition. Moreover, the objective function is concave in $\theta_H^2$. Thus, the solution to the maximization problem subject to (IIR-2$L$) is $\theta_T^{**}$ if $\theta_T^{**} \le \theta_T^*$, or $\theta_T^*$ otherwise. $\blacksquare$

*Proof of Proposition 2.* The first part of the proposition follows from Lemma 7, which provides a necessary condition for optimality for any given $\pi^1_{0d}$. Given the optimal choice of $\{\theta^t_H\}^{T+1}_{t=2}$, we may rewrite problem (3) as

$$\max_{\pi^1_{0d}} \quad qVI(\theta_H)+(1-q)(1-\theta_L)\beta\pi^1_{0d}-q\max\left\{(\theta_L-\theta_H)(\beta-G^*_T)\pi^1_{0d}-(v(\theta_H)-\theta_H\alpha),0\right\}.$$

If $v(\theta_H)-\theta_H\alpha \geq (\theta_L-\theta_H)(\beta-G^*_T)$, the objective function is strictly increasing in $\pi^1_{0d}$, and is maximized at $\pi^1_{0d}=1$. The seller needs to leave no information rent to the high type, and there is full surplus extraction. Otherwise, in Case 2(a), the objective function is still increasing in $\pi^1_{0d}$. It is optimal to set $\pi^1_{0d}=1$, but the high type will get positive information rent. In Case 2(b), the objective function is increasing on $[0,\hat{\pi}]$ and decreasing on $[\hat{\pi},1]$, where

$$\hat{\pi} := \frac{v(\theta_H)-\theta_H\alpha}{(\theta_L-\theta_H)(\beta-G^*_T)},$$

It is optimal to set $\pi^1_{0d}=\hat{\pi}$ and leave no information rent for the high type.

To complete the proof that this is the solution to problem (3), we need to show that (IIR-$ti$) holds for $i=L,H$ at all $t=2,\dots,T$. The fact that (IIR-$tH$) binds at all stages $t$ is by construction, as the stage-prices $\{p^t_L\}^T_{t=2}$ are chosen to satisfy Lemma 6. Moreover, the optimal choice of $\theta^2_H$ guarantees that (IIR-2$L$) is satisfied. We now argue that (IIR-2$L$) implies (IIR-$tL$) at $t>2$. To this end, observe that by Proposition 2, any continuation experiment sequence must be fully informative. Since the payment gap is always positive, and (IIR-3$H$) holds, (IIR-3$L$) can fail only when $VI(\theta^3_L) < VI(\theta^3_H)$, or $\theta^3_L > \sigma$. We also have shown that $\sigma > \theta^2_H$. Therefore, it suffices to prove the statement for $\theta^2_L > \theta^3_L > \sigma > \theta^2_H > \theta^3_H$. For such values of $\theta^2_L$ and $\theta^3_L$,

$$VI(\theta^3_L) = (1-\theta^3_L)\beta = \frac{(1-\theta^2_L)\pi^2_{0d}}{\phi^2_L(d)}\beta = \frac{\pi^2_{0d}}{\phi^2_L(d)}VI(\theta^2_L),$$

where $\phi^2_i(d)$ stands for the ex ante probability of obtaining a down signal in the stage experiment $e^3$ by type $i=L,H$.

Next, consider the expected payment for the mechanism. The basic recursion is

$$C(\theta^2_L;M^2_L) = p^2_L + \phi^2_L(d)C(\theta^3_L;M^3_L).$$

The expected payment for the high type also follows the same recursion, and from the fact that (IIR-$tH$) binds at $t=2,3$, we obtain

$$VI(\theta^2_H) = p^2_L + \phi^2_H(d)VI(\theta^3_H).$$

Thus,

$$\phi_L^2(d)C(\theta_L^3; M_L^3) = C(\theta_L^2; M_L^2) - \left(\theta_H^2\alpha - \phi_H^2(d)\theta_H^3\alpha\right) = C(\theta_L^2; M_L^2) - \pi_{1u}^2 VI(\theta_H^2).$$

This gives

$$\begin{aligned}\phi_L^2(d)\left[VI(\theta_L^3; E_L^3) - C(\theta_L^3; M_L^3)\right] &= \pi_{0d}^2 VI(\theta_L^2) - C(\theta_L^2; M_L^2) + \pi_{1u}^2 VI(\theta_H^2) \\ &\geq \pi_{1u}^2 VI(\theta_H^2) - (1-\pi_{0d}^2) VI(\theta_L^2),\end{aligned}$$

where the inequality comes from the fact that (IIR-2$L$) holds. Recall that we are considering the case of $VI(\theta_L^3) < VI(\theta_H^3)$ with $\theta_L^2 > \theta_L^3 > \sigma > \theta_H^2 > \theta_H^3$. Since $VI(\cdot)$ is increasing on $[0,\sigma]$ and decreasing on $[\sigma,1]$, this implies that $VI(\theta_L^2) < VI(\theta_H^2)$. Moreover, $\pi_{1u}^2 > 1-\pi_{0d}^2$ because the likelihood ratio of an up signal is greater than 1. This shows that the right-hand side of the displayed inequality is positive, and so (IIR-3$L$) holds. By applying the same argument, we can show that (IIR-$tL$) holds at any $t > 3$ if (IIR-2$L$) holds. This shows that (IIR-$ti$) holds for $i = L, H$ at all $t = 2, \dots, T$.

Finally, we have ignored (IC-$L$) in the construction. But the low type's payoff from choosing $M_L$ is 0 due to (IR-$L$), while his payoff from choosing $M_H$ would be $VI(\theta_L) - VI(\theta_H) < 0$. This ensures that (IC-$L$) also holds. ■

*Proof of Proposition 3.* Let $G_T^*(\theta_H) := G_T(\min\{\theta_T^*(\theta_H), \theta_T^{**}(\theta_H)\}; \theta_H)$, where we now explicitly write $\theta_T^*$ and $\theta_T^{**}$ as functions of $\theta_H$. Define

$$\hat{J}_T(\theta_H) := (1-\theta_H)\beta - \theta_H\alpha - (\theta_L - \theta_H)(\beta - G_T^*(\theta_H)),$$

so that Case 1 of Proposition 2 (full surplus extraction) is feasible if and only if $\hat{J}_T(\theta_H) \geq 0$. Moreover, define

$$\tilde{J}_T(\theta_H) := (1-\theta_H)\beta - \theta_H\alpha - (\theta_L - \theta_H)(\beta - G_T(\theta_T^{**}(\theta_H); \theta_H)).$$

Since $\theta_H^2 = \theta_T^{**}$ would be optimal if there were no (IIR-2$L$), we generally have $\hat{J}_T(\theta_H) \leq \tilde{J}_T(\theta_H)$. Define $\theta_T^\dagger$ to be the value that satisfies $\tilde{J}_T(\theta_T^\dagger) = 0$. Such a $\theta_T^\dagger \in (\underline{\theta}, \sigma)$ exists because $\tilde{J}_T(\underline{\theta}) > 0$ and $\tilde{J}_T(\sigma) < 0$. Also, by the envelope theorem,

$$\begin{aligned}\frac{\mathrm{d}\tilde{J}_T}{\mathrm{d}\theta_H} &= -\alpha - G_T(\theta_T^{**}; \theta_H) + (\theta_L - \theta_H)\frac{\partial G_T(\theta_T^{**}; \theta_H)}{\partial \theta_H} \\ &< -\alpha - G_T(\theta_T^{**}; \theta_H) + (\theta_L - \theta_H)\frac{1}{1-\theta_H}\alpha < 0,\end{aligned}$$

where the first inequality follows from Claim 4. It follows that $\tilde{J}_T(\cdot)$ is strictly decreasing, and thus $\theta_T^\dagger$ is unique.

Recall that $\tilde{J}_T(\theta_H) \geq 0$ ensures full surplus extraction if $\theta_H^{**}$ is indeed the optimal choice satisfying (IIR-2$L$). To see this, note first that

$$VI(\theta_L^2) = (1-\theta_L^2)\beta = \frac{(1-\theta_L)\pi_{0d}^1}{\phi_L^1(d)}\beta.$$

Note also that

$$\begin{aligned}\phi_L^1(d)C(\theta_L^2; M_L^2) = C(\theta_L; M_L) - p_L &= C(\theta_L; M_L) - \left[C(\theta_H; M_L) - \phi_H^1(d)C(\theta_H^2; M_L^2)\right] \\ &= \phi_H^1(d)VI(\theta_H^2) - \Delta C \\ &= \phi_H^1(d)\theta_H^2\alpha - \Delta C \\ &= \theta_H \pi_{0d}^1 \alpha - (\theta_L - \theta_H)\pi_{0d}^1 G(\theta_T^{**}; \theta_H).\end{aligned}$$

It then follows from these that

$$\phi_L^1(d)(VI(\theta_L^2) - C(\theta_L^2; M_L^2)) = \pi_{0d}^1\left[(1-\theta_L)\beta - \theta_H\alpha + (\theta_L - \theta_H)G(\theta_L^{**}; \theta_H)\right] = \pi_{0d}^1 \tilde{J}_T(\theta_H).$$

This shows that when $\tilde{J}_T(\theta_H) \geq 0$, (IIR-2$L$) must hold, and therefore $\theta_H^{**}$ is optimal. In other words, we have

$$\hat{J}_T(\theta_H) \begin{cases} = \tilde{J}_T(\theta_H) > 0 & \text{if } \theta_H < \theta_T^\dagger \\ = \tilde{J}_T(\theta_H) = 0 & \text{if } \theta_H = \theta_T^\dagger \\ \leq \tilde{J}_T(\theta_H) < 0 & \text{if } \theta_H > \theta_T^\dagger. \end{cases}$$

Therefore, $\theta_T^\dagger$ is also the unique crossing point of $\hat{J}_T(\cdot)$. Furthermore, since $\tilde{J}_T(\theta_H)$ decreases in $\theta_L$, by the implicit function theorem, $\theta_T^\dagger$ decreases in $\theta_L$.

When $\theta_H > \theta_T^\dagger$, Case 2 of Proposition 2 obtains. If

$$v(\theta_T^\dagger) - \theta_T^\dagger\alpha > \frac{1-q}{q}(1-\theta_L)\beta,$$

then Case 2(a) never holds, and we define $\theta_T^{\dagger\dagger} = \theta_T^\dagger$. Otherwise, we define $\theta_T^{\dagger\dagger}$ to be the value of $\theta_H$ that satisfies

$$(\theta_L - \theta_H)(\beta - G_T^*) - \frac{1-q}{q}(1-\theta_L)\beta = 0. \tag{16}$$

Note that

$$\frac{\mathrm{d}G_T^*}{\mathrm{d}\theta_H} = \begin{cases} \dfrac{\partial G_T(\theta_T^{**};\theta_H)}{\partial \theta_H} > 0 & \text{if } \theta_T^{**} < \theta_T^* \\ \dfrac{\partial G_T(\theta_T^*;\theta_H)}{\partial \theta_H^2}\dfrac{\mathrm{d}\theta_T^*}{\mathrm{d}\theta_H} + \dfrac{\partial G_T(\theta_T^*;\theta_H)}{\partial \theta_H} > 0 & \text{if } \theta_T^{**} \geq \theta_T^*, \end{cases}$$

The second part of the above follows because $\partial G_T/\partial \theta_H^2 > 0$ at any $\theta_H^2 < \theta_T^{**}$, and because $\mathrm{d}\theta_T^*/\mathrm{d}\theta_H > 0$ (which follows from implicit differentiation of $J_T^2(\theta_T^*;\theta_H) = 0$). This shows that the left-hand side of (16) is strictly decreasing in $\theta_H$, and hence $\theta_T^{\dagger\dagger}$ is unique. Moreover, Case 2(a) obtains if $\theta_H \in (\theta_T^\dagger, \theta_T^{\dagger\dagger})$, while Case 2(b) obtains if $\theta_H \geq \theta_T^{\dagger\dagger}$. ■

*Proof of Proposition 6.* Since $\theta_T^{**} = 1-(1-\theta_H)^{\frac{T-1}{T}}$, it is clear that $\lim_{T\to\infty} \theta_T^{**} = \theta_H$. The limiting value $\theta^*$ is the value of $\theta_H^2$ that solves:

$$J_\infty^2(\theta_H^2;\theta_H) := (1-\theta_H^2)\beta - \theta_H^2\alpha - (\rho(\theta_H^2;\theta_H) - \theta_H^2)\big(\beta - G_\infty^2(\theta_H^2)\big) = 0,$$

where, similar to the derivation of $G_\infty(\theta_H^2;\theta_H)$,

$$G_\infty^2(\theta_H^2) = -\frac{\alpha}{\theta_H^2}\big(\theta_H^2 + \ln(1-\theta_H^2)\big) = G_\infty(\theta_H^2;\theta_H^2).$$

Hence, $\theta^*$ is given by equation (12).

Now, define

$$J_\infty(\theta_H^2;\theta_H) := (1-\theta_H)\beta - \theta_H\alpha - (\theta_L - \theta_H)\big(\beta - G_\infty(\theta_H^2;\theta_H)\big).$$

Then, $\theta^\dagger$ given by equation (11) satisfies $J_\infty(\theta^\dagger;\theta^\dagger) = 0$. Furthermore, at $\theta_H^2 = \theta^\dagger$ and $\theta_H = \theta^\dagger$, we have

$$J_\infty^2(\theta_H^2;\theta_H) = (1-\theta^\dagger)\beta - \theta^\dagger\alpha - (\rho(\theta^\dagger;\theta^\dagger) - \theta^\dagger)\big(\beta - G_\infty^2(\theta^\dagger)\big) = J_\infty(\theta^\dagger;\theta^\dagger) = 0,$$

where the last equality follows because $\rho(\theta_H;\theta_H) = \theta_L$ and $G_\infty^2(\theta_H) = G_\infty(\theta_H;\theta_H)$ for any $\theta_H$. This shows that $\theta^*$ is equal to $\theta^\dagger$ when $\theta_H = \theta^\dagger$. Moreover, because $\theta^*$ increases in $\theta_H$, we have $\theta^* \geq \theta_H$ if $\theta_H < \theta^\dagger$ and $\theta^* \geq \theta_H$ if $\theta_H \geq \theta^\dagger$.

Given this, if $\theta_H < \theta^\dagger$, then (IIR-2$L$) is slack at $\theta_H^2 = \min\{\theta_H, \theta^*\} = \theta_H$. We have $J_\infty(\theta_H;\theta_H) > J_\infty(\theta^\dagger;\theta^\dagger) = 0$. This implies that full surplus extraction is achievable in the limit mechanism. For finite mechanisms, the fact that $\theta_T^\dagger$ increases monotonically with $T$ with a limiting value of $\theta^\dagger$ implies that there exists a sufficiently large and finite number of stages $T'$ such that $\theta_H \leq \theta_{T'}^\dagger$ if and only if $\theta_H < \theta^\dagger$. By Proposition 3, it is optimal to choose such a $T'$ so that full surplus extraction is achieved. ■

## References

**Alonso, Ricardo and Odilon Camara**, "Persuading Voters," *American Economic Review*, 2016, *106* (11), 3590–3605.

**Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin**, "The Design and Price of Information," *American Economic Review*, 2018, *108* (1), 1–48.

_ **and** _ , "Markets for Information: An Introduction," *Annual Review of Economics*, 2019, *11* (1), 85–107.

_ **, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter**, "Optimal Information Disclosure in Classic Auctions," *American Economic Review: Insights*, 2022, *4* (3), 371–388.

**Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang**, "Pivotal Persuasion," *Journal of Economic Theory*, 2019, *180*, 178–202.

**Coval, Joshua D. and Anjan V. Thakor**, "Financial Intermediation as a Beliefs-Bridge between Optimists and Pessimists," *Journal of Financial Economics*, 2005, *75* (3), 535–569.

**Crémer, Jacques and Richard P. McLean**, "''Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 1985, *53* (2), 345–361.

_ **and** _ , "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1988, *56* (6), 1247–1257.

**Eső, Peter and Balazs Szentes**, "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 2007, *74* (3), 705–731.

**Guo, Yingni and Eran Shmaya**, "The Interval Structure of Optimal Disclosure," *Econometrica*, 2019, *87* (2), 653–675.

**Haghpanah, Nima and Ron Siegel**, "Screening Two Types," *Working Paper, Pennsylvania State University*, 2025.

**Hörner, Johannes and Andrzej Skrzypacz**, "Selling Information," *Journal of Political Economy*, 2016, *124* (4), 1515–1562.

**Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li**, "Persuasion of a Privately Informed Receiver," *Econometrica*, 2017, *85* (6), 1949–1964.

**Krähmer, Daniel**, "Information Disclosure and Full Surplus Extraction in Mechanism Design," *Journal of Economic Theory*, 2020, *187*, 105020.

**Li, Hao and Xianwen Shi**, "Discriminatory Information Disclosure," *American Economic Review*, 2017, *107* (11), 3363–3385.

**Liu, Heng**, "Efficient Dynamic Mechanisms in Environments with Interdependent Valuations: the Role of Contingent Transfers," *Theoretical Economics*, 2018, *13* (2), 795–830.

**Maskin, Eric and John Riley**, "Monopoly with Incomplete Information," *RAND Journal of Economics*, 1984, *15* (2), 171–196.

**McAfee, R. Preston and Philip J. Reny**, "Correlated Information and Mechanism Design," *Econometrica*, 1992, *60* (2), 395–421.

**Morris, Stephen**, "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, November 1994, *62* (6), 1327–1347.

**Mussa, Michael and Sherwin Rosen**, "Monopoly and Product Quality," *Journal of Economic Theory*, 1978, *18* (2), 301–317.

**Noda, Shunya**, "Full Surplus Extraction and Within-period ex post Implementation in Dynamic Environments," *Theoretical Economics*, 2019, *14* (1), 39–69.

**Sebenius, James K. and John Geanakoplos**, "Don't Bet on it: Contingent Agreements with Asymmetric Information," *Journal of the American Statistical Association*, 1983, *78* (382), 424–426.

**Wei, Dong and Brett Green**, "(Reverse) Price Discrimination with Information Design," *American Economic Journal: Microeconomics*, 2024, *16* (2), 267–295.