Selling Information Sequentially

Vitalii Tubdenov*

Abstract

A seller sequentially supplies information to a decision-maker under uncertainty. In each period, the seller posts a price for the current signal, and the buyer either stops and acts or buys the signal. After each sale, the signal realization is observed by both parties. I characterize equilibrium under a common prior. In a Nash equilibrium, the seller extracts the full social surplus using a two-part tariff with an access fee and per-period prices set at marginal cost. In subgame-perfect equilibria, full extraction persists with history-dependent pricing, which may involve prices below marginal cost.

JEL: D83, D81

Keywords: selling information, dynamic pricing, sequential learning, optimal stopping

Economic agents do not gather the facts needed to make a decision at once. They learn in stages, and after each new piece of information, they decide whether to act immediately (and how) or to keep investigating. In many settings, such information is produced and sold by specialist vendors. For example, pharmaceutical companies hire contract research organizations (CROs) to run phases of clinical research and then, after each phase, decide whether to continue or discontinue drug development.[1]

In this paper, I study selling information in dynamic settings. The current literature on selling information focuses mainly on static problems in which the seller designs the selling mechanism, possibly controlling the information being sold. In dynamic settings, each new piece of evidence updates the buyer's beliefs, which changes the buyer's willingness to pay for future information. As the buyer's beliefs change during the interaction, how should the seller dynamically price information?

*Tubdenov: Department of Economics, Harvard University (email: vtubdenov@g.harvard.edu). I am grateful to Shengwu Li, Jesse Shapiro, Rahul Singh, and Elie Tamer for guidance. I thank participants in the graduate student workshops at Harvard and the TTIC workshop on Incentives for Collaborative Learning and Data Sharing for comments.

[1]Pharmaceutical companies discontinue projects when products fail safety or efficacy requirements or produce undesired clinical results (ICON plc 2025). Depending on the phase, 25% to 70% of drugs advance to the next phase (FDA 2018).

To study this question, I model the buyer's problem as a sequential statistical decision problem (Wald 1945). The buyer faces a single terminal decision under uncertainty and therefore values learning before acting. He observes evidence one piece at a time and, after each observation, decides whether to stop and take an action or to continue sampling and act later. In statistics, the buyer's cost of acquiring observations is taken as exogenous. In practice, however, the buyer often cannot gather evidence on his own and instead relies on information vendors. I incorporate such vendors into the sequential decision problem.

As an illustration, consider a common arrangement: a CRO contracts with a pharmaceutical company to conduct phases of clinical research. Such relationships employ a variety of payment schemes. See, e.g., Birks and Singh (2011); PPD (2026). A common baseline is input-based pricing. An alternative is a predetermined, all-inclusive flat rate for a defined scope of services. Importantly, outcome-based contracts, which tie the fee to the clinical program's success, are also used.

In my model, a monopolist data seller provides information through a finite sequence of signals, and producing each signal entails a deterministic marginal cost that may vary across periods. In each period, the seller posts a price for the current signal. The buyer either stops and takes the terminal action or pays to observe the next signal. The joint distribution of the state and the full signal process is common knowledge, and after each purchase, the signal realization is observed by both parties. Hence, neither party has private information about the state or about realized signals at any time.

My first set of results concerns Nash equilibrium (NE) of the game. I construct a seller-preferred NE with a two-part tariff (Oi 1971). First, the seller charges a fixed access fee, which depends on the buyer's prior beliefs and the decision problem, as well as the signals being sold. Second, per-period prices are set at marginal cost and are therefore history-independent.

The two-part tariff allows the seller to extract the efficient social surplus. Setting per-period prices at marginal cost makes the buyer face the efficient marginal incentive to continue sampling, so the buyer's stopping rule is efficient. The fixed access fee is then set equal to the value of efficient learning, making the buyer indifferent between participating in the interaction and acting without acquiring additional evidence whenever buying the first signal is efficient.

Given full extraction, a key object is the social surplus — the value of sequential learning under efficient information acquisition. I show that the social planner's problem reduces to a collection of static terminal decision problems together with an optimal stopping problem. This decomposition yields a simple characterization of the efficient stopping and terminal decision rules, and of the social surplus.

My next set of results concerns subgame-perfect equilibrium (SPE) of the game. In that case, both players' strategies must be sequentially rational after every history of past signal realizations and prices.

The two-part tariff is generally not sequentially rational: after the access fee has been paid, the seller may prefer to raise continuation prices above marginal cost. The constructed Nash equilibrium deters such deviations through the buyer's rejection of off-schedule prices, which need not be sequentially rational. In subgame-perfect equilibrium, prices are generally history-dependent because the buyer's willingness to pay can vary with realized evidence.

I characterize equilibrium pricing recursively and show that full social surplus extraction nevertheless persists in every SPE. The key forces are that the parties begin from a common prior, and that each realized signal is publicly observed. This makes the buyer's continuation value common knowledge. The seller can then set the price so that the buyer is just willing to continue in histories where continuing is socially efficient, and to stop in histories where stopping is socially efficient.

The resulting pricing policies may be complex. In particular, they may involve offering signals at prices below marginal cost. This occurs, for example, when the first signal has little standalone value for the buyer regardless of its realization, while selling it makes profitable later sales possible.

More broadly, the paper is about selling information throughout a learning process rather than the sale of a single piece of information. This perspective highlights how the timing of information acquisition shapes the value of efficient learning. The analysis also highlights that there are two distinct economic tasks in these environments: implementing efficient learning and extracting the value created by efficient learning.

The rest of the paper proceeds as follows. Section 1 reviews the related literature. Section 2 introduces the model. Section 3 presents the results. Section 4 concludes.

1 Related Literature

I contribute to the literature on selling information. For a review, see Bergemann and Bonatti (2019). Since Blackwell's foundational contributions (Blackwell 1951, 1953), there has been growing interest in understanding how the value generated by information is traded in markets, beginning with the seminal work of Admati and Pfleiderer (1986, 1990). My main contribution is to consider a setting in which the value of information comes not only from informing which decision to take, but also when to take it.

Lehrer and Wang (2024) study selling information in a sequential decision-making framework. In their model, signals are conditionally independent and costless to produce, so an

upfront fee extracts the efficient social surplus. I generalize their result by allowing for signals with time-dependent costs and a general joint distribution. Further, I show that full surplus extraction persists in SPE, and characterize an optimal pricing scheme.

Chen, Ishida and Suen (2026) study a monopolist designing a menu consisting of sequences of signals with pricing schedules to screen buyers with privately known initial beliefs.[2] They show that, in a binary-state, binary-action environment, the seller can benefit from offering sequences rather than single signals. I abstract from screening and signal design: the sequence of experiments is exogenously given and beliefs are common knowledge, shifting attention to the pricing problem. Further, I allow for general state spaces and decision problems, under NE and SPE.

A number of contributions focus on specialized dynamic settings. Hörner and Skrzypacz (2016) show how gradual information disclosure can mitigate the holdup problem faced by an agent proving her competence to a firm. Mekonnen, Murra-Anton and Pakzad-Hurson (2025) study sequential search by an agent who can purchase signals about the quality of options from a seller with limited commitment. Rivera Mora and Rodríguez (2026) study information selling under prior disagreement and limited commitment.

I also contribute to the sequential sampling literature by endogenizing the cost of each sample and modeling it as the outcome of a pricing problem, building on Wald (1945). Morris and Strack (2019) link sequential sampling to ex ante information costs. Further, Kruse and Strack (2015) study how an interested party can influence the stopping decision using transfers, but without the ability to observe the decision-maker's signals.

2 Model

There are two agents, a buyer (he) and a seller (she). Time is discrete and finite, $t = 1, \dots, T$. The buyer makes a terminal decision but is uncertain about a payoff-relevant state. The seller supplies signals to the buyer, one in each period. Selling signal t entails marginal cost c_t for the seller. The joint distribution of the state and the signal process is common knowledge.

In each period, the seller sets the price for the current signal. The buyer either stops and takes his decision, and the interaction ends, or purchases the period's signal at the posted price, and the interaction continues. Upon purchase, the signal realization is observed by both parties. Hence, neither party has private information about the state or signal realizations in any period. See Appendix A for technical conditions.

[2] Chen, Ishida and Suen (2026) build on Bergemann, Bonatti and Smolin (2018), who study the design and pricing of a menu of experiments. Other contributions to designing and selling information in static settings include Liu, Shen and Xu (2021); Eső and Szentes (2007); Bonatti, Dahleh, Horel and Nouripour (2024); Rodríguez Olivera (2024).

2.1 Decision problem

Given the unknown state $\theta \in \Theta$, the buyer chooses an action $a \in \mathcal{A}$ to minimize the loss function $L : \Theta \times \mathcal{A} \to [0, \bar{L}]$, for some finite $\bar{L}$. With an abuse of notation, for $\sigma \in \Delta(\mathcal{A})$, write $L(\theta, \sigma) \coloneqq \int_{\mathcal{A}} L(\theta, a)\sigma(da)$.

2.2 Signal process

The period-t signal is a random variable $X_t \in \mathcal{X}_t$, with realization x_t. Let $\mathbf{H}_t \coloneqq (X_1, \ldots, X_t)$ denote the signal history through period t, and let $\mathbf{h}_t \coloneqq (x_1, \ldots, x_t)$ denote a realization, with $\mathbf{h}_0 \coloneqq \varnothing$. For $t \geq 1$, write $\mathcal{X}^t \coloneqq \prod_{s=1}^{t} \mathcal{X}_s$, and let $\mathcal{X}^0$ be a singleton. Let μ denote the joint distribution of $(\theta, X_1, \ldots, X_T)$, and π_0 be the marginal distribution of θ, and let $\pi_t(\cdot \mid \mathbf{H}_t)$ be a conditional distribution of θ given $\mathbf{H}_t$.

2.3 Timing

In each period t before the game ends, the seller posts a price p_t for the signal X_t, and the buyer either stops and takes an action or purchases the signal. If he purchases, both parties observe the realization $X_t = x_t$, and the game continues to period $t+1$. Let $\tau \in \{0, 1, \ldots, T\}$ denote the number of signals purchased.

2.4 Strategies

The buyer's strategy consists of a sequence of stopping rules $\varphi = \{\varphi_j\}_{j=0}^{T}$ and a sequence of terminal decision rules $\delta = \{\delta_j\}_{j=0}^{T}$. A stopping rule φ_j prescribes the probability of stopping after observing the signal realizations up to period j (observing no signals for $j = 0$) and having been quoted p_{j+1}. Let $\mathbf{p}_{j+1} \coloneqq (p_1, \ldots, p_{j+1})$ denote the price history through period $j + 1$, with $\mathbf{p}_0 \coloneqq \varnothing$. Then, fixing $\varphi_T \equiv 1$,

$$\varphi = \{\varphi_0(\mathbf{h}_0, \mathbf{p}_1), \varphi_1(\mathbf{h}_1, \mathbf{p}_2), \ldots, \varphi_{T-1}(\mathbf{h}_{T-1}, \mathbf{p}_T), \varphi_T\}.$$

A terminal decision rule δ_j prescribes a distribution over actions if the buyer stops after observing the signal realizations up to period j (observing no signals for $j = 0$) and having been quoted p_{j+1} (without purchasing). Thus,

$$\delta = \{\delta_0(\mathbf{h}_0, \mathbf{p}_1), \delta_1(\mathbf{h}_1, \mathbf{p}_2), \ldots, \delta_{T-1}(\mathbf{h}_{T-1}, \mathbf{p}_T), \delta_T(\mathbf{h}_T, \mathbf{p}_T)\}.$$

The seller's strategy specifies which prices to quote: $p_t(\mathbf{h}_{t-1}, \mathbf{p}_{t-1}) \in \mathbb{R}$.

2.5 Payoffs

The buyer's payoff is $u_B \coloneqq -L(\theta, a) - \sum_{t=1}^{\tau} p_t$. The seller faces a marginal cost $c_t \in \mathbb{R}$ of providing signal t, known in advance to both parties. The seller's payoff is revenue net of production costs: $u_S \coloneqq \sum_{t=1}^{\tau}(p_t - c_t)$.

2.6 Solution

I consider NE and SPE using an extensive form with decision nodes indexed by signal and price histories. Nature draws each signal according to its distribution conditional on past signals. If the buyer stops after $\mathbf{h}_t$ and chooses a, his terminal loss is $\int_\Theta L(\theta, a)\, \pi_t(d\theta \mid \mathbf{h}_t)$. Every decision node starts a subgame.

3 Results

3.1 Efficient Allocation

A social planner chooses a stopping rule φ and a terminal decision rule δ to maximize the expected sum of the agents' payoffs. Equivalently, the social planner chooses (φ, δ) to minimize the Bayes risk $r(\varphi, \delta)$:

$$r(\varphi, \delta) \coloneqq \mathrm{E}\left(\sum_{t=0}^{T} \psi_t(\mathbf{H}_t)[L\{\theta, \delta_t(\mathbf{H}_t)\} + \sum_{s=1}^{t} c_s]\right), \tag{1}$$

where $\psi_0(\mathbf{h}_0) \coloneqq \varphi_0(\mathbf{h}_0)$, and for $t \geq 1$, $\psi_t(\mathbf{h}_t) \coloneqq \varphi_t(\mathbf{h}_t) \prod_{s=0}^{t-1}\{1 - \varphi_s(\mathbf{h}_s)\}$ is the conditional probability that the planner stops at t. The efficient rules depend only on the signal history because the prices cancel out.

Proposition 1 shows that, holding the stopping rule fixed, the planner's terminal decision can be chosen period by period and history by history. Given these terminal decisions, Proposition 2 shows that the optimal stopping rule compares the value of stopping immediately to the value of continuing one more period and behaving optimally thereafter.

I impose the following regularity condition.

Assumption 1. *For each $t = 0, \ldots, T$, there exists a measurable function $\delta_t^0 : \mathcal{X}^t \to \Delta(\mathcal{A})$ such that, for every history $\mathbf{h}_t$,*

$$\mathrm{E}\left[L\left\{\theta, \delta_t^0(\mathbf{h}_t)\right\} \mid \mathbf{H}_t = \mathbf{h}_t\right] = \inf_{\sigma \in \Delta(\mathcal{A})} \mathrm{E}\left\{L(\theta, \sigma) \mid \mathbf{H}_t = \mathbf{h}_t\right\}.$$

I define the posterior Bayes risk by

$$\rho_t(\mathbf{h}_t) \coloneqq \mathrm{E}\left[L\left\{\theta, \delta_t^0(\mathbf{h}_t)\right\} \mid \mathbf{H}_t = \mathbf{h}_t\right].$$

A sufficient condition for Assumption 1 is that $\mathcal{A}$ is finite.

I start with the first step. Fix a stopping rule. An optimal terminal decision can then be chosen period by period and history by history: for each period t and each realized history $\mathbf{h}_t$, the planner selects an optimal terminal action for the induced static decision problem. The following proposition formalizes this insight. All proofs are in Appendix B.

Proposition 1 (Efficient terminal decision rule)**.** *For measurable functions δ_t^0, for $t = 0, \ldots, T$, satisfying Assumption 1, write $\delta^0 = (\delta_t^0)_{t=0}^T$. Then, for every stopping rule φ and terminal decision rule δ, $r(\varphi, \delta^0) \leq r(\varphi, \delta)$.*

Conditional on stopping at time t, the planner faces a static decision problem and selects a Bayes rule. For a fixed stopping rule, a terminal decision rule is optimal if and only if it is a Bayes rule at almost every stopping history.

I turn to the problem of finding an optimal stopping rule. Fix δ^0 as in Proposition 1, and define the social value function recursively. Set $V_T(\mathbf{h}_T) \coloneqq \rho_T(\mathbf{h}_T) + \sum_{s=1}^T c_s$, and for $t < T$ define

$$V_t(\mathbf{h}_t) \coloneqq \min\left[\rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s,\ \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}\right].$$

Proposition 2 (Efficient stopping rule)**.** *Consider φ^0 such that for each $t < T$ and each history $\mathbf{h}_t$,*

$$\varphi_t^0(\mathbf{h}_t) = \begin{cases} 1, & \text{if } \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s < \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}, \\ \in [0,1], & \text{if } \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s = \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}, \\ 0, & \text{if } \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s > \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}. \end{cases}$$

Then (φ^0, δ^0) solves the social planner's problem.

At each history, the planner compares the loss from stopping immediately with the expected loss from continuing one more period and behaving optimally thereafter. I now define the social surplus.

Definition 1 (Social surplus)**.** *For all $t \in \{0, \ldots, T\}$ and all histories $\mathbf{h}_t$, the continuation social surplus relative to stopping is $S_t(\mathbf{h}_t) \coloneqq \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s - V_t(\mathbf{h}_t)$. I set $S = S_0(\mathbf{h}_0)$ and refer to S as the social surplus.*

This definition makes clear the determinants of the social surplus. At history $\mathbf{h}_t$, the term $\rho_t(\mathbf{h}_t)$ is the expected loss if the planner stops immediately and takes the best action given current beliefs. The value $V_t(\mathbf{h}_t)$ instead incorporates the option to continue learning, taking into account the marginal sampling costs and the conditional distribution of future beliefs. Therefore, $S_t(\mathbf{h}_t)$ measures the value of the remaining learning opportunity generated by the model fundamentals: the loss function, the buyer's current posterior, the informativeness of future signals, and the cost of acquiring them.

3.2 Nash Equilibrium

The seller can achieve full extraction with a two-part tariff (Oi 1971), separating implementation from extraction. She first sets per-period prices to align the buyer's incentive to purchase the next signal at each history with the efficient stopping behavior. She then extracts the full social surplus through an access fee, which does not distort any continuation decision.

Theorem 1 (Surplus extraction in NE)**.** *Let (φ^0, δ^0) be a socially efficient allocation from Propositions 1 and 2, such that*

$$\varphi_t^0(\mathbf{h}_t) = \begin{cases} 1, & \text{if } \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s < \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}, \\ 0, & \text{if } \rho_t(\mathbf{h}_t) + \sum_{s=1}^t c_s \geq \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}. \end{cases}$$

Consider the following strategy profile.
Seller. *In each period $t = 1, \ldots, T$, the seller posts price p_t^*:*

$$p_1^* \coloneqq c_1 + S, \qquad p_t^* \coloneqq c_t \quad \text{for all } t = 2, \ldots, T.$$

Buyer. *For each $t = 0, 1, \ldots, T-1$ and each history $(\mathbf{h}_t, \mathbf{p}_{t+1})$, set*

$$\{\varphi_t^*(\mathbf{h}_t, \mathbf{p}_{t+1}), \delta_t^*(\mathbf{h}_t, \mathbf{p}_{t+1})\} \coloneqq \begin{cases} \{\varphi_t^0(\mathbf{h}_t), \delta_t^0(\mathbf{h}_t)\}, & \text{if } \mathbf{p}_{t+1} = (p_1^*, \ldots, p_{t+1}^*), \\ \{1, \delta_t^0(\mathbf{h}_t)\}, & \text{otherwise,} \end{cases}$$

and set $\varphi_T^* \equiv 1$ and $\delta_T^*(\mathbf{h}_T, \mathbf{p}_T) \coloneqq \delta_T^0(\mathbf{h}_T)$.

Then $\{(\varphi^*, \delta^*), p^*\}$ is an NE. Moreover, along the equilibrium path the buyer follows the socially efficient allocation (φ^0, δ^0). The seller's realized equilibrium payoff equals the social surplus S and the buyer's expected equilibrium payoff equals $-\rho_0$, the buyer's maximal expected payoff without learning.

The buyer faces marginal cost for every signal after the first. As a result, his continuation decisions coincide with the planner's. The equilibrium therefore implements the socially efficient allocation (φ^0, δ^0).

The role of the first-period price is extraction. Before the access fee, at per-period prices equal to c_t, the buyer's expected gain from entering the relationship relative to stopping immediately equals the social surplus S whenever buying the first signal is efficient; otherwise $S = 0$ and the buyer stops immediately. The seller therefore sets $p_1^* = c_1 + S$: the component c_1 covers the cost of the first signal, and S is a fixed access fee. Assuming that the buyer purchases when indifferent, the seller's payoff equals S.

The NE in Theorem 1 is seller-preferred. It is not unique in a substantive sense: there are NE with no trade. Fix any price path $\tilde{p} = (\tilde{p}_1, \ldots, \tilde{p}_T) > \mathbf{0}$ such that $\tilde{p}_1 > \sup_{(\theta,a)} L(\theta, a)$, so the buyer strictly prefers stopping immediately to purchasing in period 1. Consider the strategy profile in which the seller posts $\tilde{p}_t$ in each period t, and the buyer stops at $t = 0$ with probability one, choosing a Bayes rule: $\varphi_0 \equiv 1$ and $\delta_0 \in \arg\min_{\sigma \in \Delta(\mathcal{A})} \mathrm{E}\{L(\theta, \sigma)\}$. This profile is an NE: given $\tilde{p}_1 > \sup_{(\theta,a)} L(\theta, a)$, stopping immediately is part of a best response for the buyer, and given that the buyer stops immediately regardless of prices, the seller is indifferent among all pricing strategies and in particular is best responding by posting $\tilde{p}$.

3.3 Subgame-Perfect Equilibrium

SPE imposes sequential rationality: after every history, consisting of the signal history and the price history, both players' continuation strategies must be optimal in the corresponding subgame. NE relies on the buyer's rejection of off-schedule prices, which need not be sequentially rational. Once the access fee has been paid, the seller may prefer to raise continuation prices above marginal cost to extract surplus left to the buyer if he responds optimally to the new price. Thus, the two-part tariff need not be sequentially rational.

The analysis therefore asks whether full social surplus extraction survives the requirement of sequential rationality, and how equilibrium prices adjust over time as evidence arrives.

Full social surplus extraction holds in every SPE: because beliefs and payoffs are common knowledge at every history, and the seller makes take-it-or-leave-it offers, the buyer earns no informational rents in equilibrium. The key difference from NE is the pricing strategy. In SPE, prices are generally history-dependent because realized evidence can change the buyer's willingness to pay for additional signals. I characterize SPE recursively, yielding a simple expression for the seller's optimal price after each history.

In what follows, continuation payoffs exclude past payments and costs and, for $t < T$, are measured before the next quote.

Theorem 2 (Surplus extraction in SPE)**.** *Let (φ^0, δ^0) be a socially efficient allocation from Propositions 1 and 2. Consider the following strategy profile.*
Seller. *For each $t \in \{0, 1, \ldots, T-1\}$ and each history $(\mathbf{h}_t, \mathbf{p}_t)$, set*

$$p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t, \mathbf{p}_t) = \begin{cases} \rho_t(\mathbf{h}_t) - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}, & \text{if } S_t(\mathbf{h}_t) > 0, \\ \rho_t(\mathbf{h}_t) - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\} + 1, & \text{if } S_t(\mathbf{h}_t) = 0. \end{cases}$$

Buyer. *For each $t < T$ and each history $(\mathbf{h}_t, \mathbf{p}_{t+1})$, set $\delta_t^{\mathrm{SPE}}(\mathbf{h}_t, \mathbf{p}_{t+1}) = \delta_t^0(\mathbf{h}_t)$. Let the buyer purchase signal $t+1$ if and only if $p_{t+1} \leq \rho_t(\mathbf{h}_t) - \mathrm{E}\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}$. Finally, set $\varphi_T^{\mathrm{SPE}} \equiv 1$ and $\delta_T^{\mathrm{SPE}}(\mathbf{h}_T, \mathbf{p}_T) = \delta_T^0(\mathbf{h}_T)$.*

Then $\left\{(\varphi^{\mathrm{SPE}}, \delta^{\mathrm{SPE}}), p^{\mathrm{SPE}}\right\}$ is an SPE. Moreover, for every t, $\mathbf{h}_t$, and $\mathbf{p}_t$, the seller's continuation payoff equals $S_t(\mathbf{h}_t)$ and the buyer's continuation payoff equals $-\rho_t(\mathbf{h}_t)$. These payoffs are the same in every SPE.

Along the equilibrium path, the buyer's stopping rule is efficient. The seller's expected payoff equals the social surplus S, while the buyer's expected payoff equals $-\rho_0$. Compared to the constructed NE, where the seller obtains S for sure, the seller in SPE gets S in expectation, but her realized payoff may vary with the signal history.

Although the equilibrium payoff characterization is simple, the associated pricing scheme need not be. Sequential rationality requires prices to adjust to the realized signal history, because new realizations can change the buyer's continuation value and hence his willingness to pay for future information. As a result, equilibrium prices are generally history-dependent and can vary substantially across continuation subgames, even though the induced allocation remains efficient and the seller continues to extract the full social surplus in expectation.

This recursive characterization also shows that equilibrium prices need not be bounded below by marginal cost. In some histories, the seller may price the current signal below marginal cost in order to move the game to a subgame in which future signals are substantially more valuable.

On histories where trade occurs ($S_t(\mathbf{h}_t) > 0$), the equilibrium price can be rewritten as $p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t) = c_{t+1} + S_t(\mathbf{h}_t) - \mathrm{E}\left\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}$. Therefore $p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t) < c_{t+1}$ if and only if $\mathrm{E}\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} > S_t(\mathbf{h}_t)$, i.e., when the expected continuation surplus increases after observing one more signal.

The next example illustrates this mechanism in a simple environment. The first signal is only mildly informative, while the second signal is perfectly revealing. As a result, the seller is willing to price the first signal below marginal cost in order to reach the more profitable continuation subgame.

Example 1 (Below-cost pricing). *Let* $T = 2$, $\Theta = \mathcal{A} = \{0, 1\}$, *and* $L(\theta, a) = 1\{\theta \neq a\}$. *Assume* $\Pr(\theta = 1) = \Pr(\theta = 0) = 0.5$, *and* $c_1 = c_2 = 0.2$. *The first signal* $X_1 \in \{g, b\}$ *is mildly informative:* $\Pr(g \mid \theta = 1) = \Pr(b \mid \theta = 0) = 0.6$. *The second signal is perfectly revealing:* $X_2 = \theta$. *Then at* $t = 0$, *the posterior Bayes risk is* $\rho_0 = 0.5$. *After either realization of* X_1, $\rho_1(\mathbf{h}_1) = 0.4$. *After observing* X_2, *the state is known exactly, so* $\rho_2(\mathbf{h}_2) = 0$.

After the first signal is purchased, the seller charges $p_2^{\mathrm{SPE}}(\mathbf{h}_1) = \rho_1(\mathbf{h}_1) - \mathrm{E}\{\rho_2(\mathbf{H}_2) \mid \mathbf{H}_1 = \mathbf{h}_1\} = 0.4$, *so the continuation payoff is* $S_1(\mathbf{h}_1) = 0.2$.

At the initial node, $p_1^{\mathrm{SPE}} = \rho_0 - \mathrm{E}\{\rho_1(\mathbf{H}_1)\} = 0.1 < c_1$. *Thus the first signal is sold below marginal cost.*

The reason is that the first signal has little standalone value: it reduces Bayes risk only from 0.5 *to* 0.4. *But selling it moves the game to period* 2, *where the seller earns continuation payoff* 0.2. *The seller is willing to price the first signal below cost, and her total equilibrium payoff is* $S_0 = 0.1$.

4 Discussion

This paper studies selling information when a buyer learns sequentially and decides when to stop. I show that the pricing problem involves two tasks: implementing efficient learning and extracting the value it creates. In an NE, a two-part tariff separates these tasks: per-period prices equal to marginal cost implement efficient stopping, while a fixed fee extracts the surplus. In SPE, pricing generally becomes history-dependent because the buyer's willingness to pay changes with realized evidence. Even so, full surplus extraction survives in every SPE. The buyer's continuation values are known by the seller after every history, allowing the seller to adjust prices without distorting efficient stopping.

The results rely on strong assumptions. I briefly discuss possible future developments.

The first would be to introduce private information on the buyer's side, such as a privately known loss function. Then the seller no longer knows the buyer's continuation value after each history, so dynamic pricing becomes a screening problem, and full surplus extraction generally fails.

The second is to allow for information design. In my model, the sequence of experiments is exogenously given, so the seller only chooses prices. Chen, Ishida and Suen (2026) make progress in this direction by incorporating both buyer private information and signal design in a simple environment with a binary state and binary action.

The third direction is to consider settings beyond bilateral monopoly. With multiple sellers, a central question is how complementarities or substitutabilities across sellers' information shape prices and the buyer's learning. With multiple buyers, sellers may want

to design contracts that reflect the fact that one buyer's purchase can change the value of information for others—for example, when the buyers are competing firms in downstream markets.

A Technical Conditions

Θ, $\mathcal{A}$ and $\mathcal{X}_t$ are Polish spaces, equipped with Borel σ-algebras. Equip $\Delta(\mathcal{A})$ with the Borel σ-algebra induced by the weak topology. All finite products carry the product σ-algebra. The loss function L is measurable.

Fix versions under μ of the regular conditional laws of X_{t+1} given $\mathbf{H}_t$, $t = 0, \ldots, T-1$, and of the terminal posterior π_T. These laws are common knowledge at every history. Conditional expectations given a signal history use the continuation laws obtained by iterating the chosen next-signal laws and applying π_T. Define earlier posteriors backward by

$$\pi_t(A \mid \mathbf{h}_t) = \mathrm{E}\{\pi_{t+1}(A \mid \mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}$$

for $t = T-1, \ldots, 0$, every history $\mathbf{h}_t$ and every Borel set $A \subseteq \Theta$. The resulting π_t are versions under μ of the conditional laws of θ given $\mathbf{H}_t$, with consistency across dates holding at every history, including histories of probability zero.

All strategy components are Borel measurable on the signal and price history spaces specified in Section 2.

Expected payoffs may take extended-real values and must be well defined ex ante for NE and in every subgame for SPE. Equilibrium comparisons are restricted to unilateral deviations whose expected payoff is well defined in the game or subgame under consideration.

B Proofs

Given a strategy profile $\{(\varphi, \delta), p\}$, let $\mathbf{P}_t = (P_1, \ldots, P_t)$ denote the random price history induced by p, where $\mathbf{P}_0 \coloneqq \varnothing$, $P_1 \coloneqq p_1(\mathbf{H}_0, \mathbf{P}_0)$, and $P_t \coloneqq p_t(\mathbf{H}_{t-1}, \mathbf{P}_{t-1})$ for $t = 2, \ldots, T$. Define the induced probability of first stopping at each date by $\psi_0 \coloneqq \varphi_0(\mathbf{H}_0, \mathbf{P}_1)$,

$$\psi_t \coloneqq \varphi_t(\mathbf{H}_t, \mathbf{P}_{t+1}) \prod_{s=0}^{t-1} \{1 - \varphi_s(\mathbf{H}_s, \mathbf{P}_{s+1})\}, \qquad t = 1, \ldots, T-1,$$

and $\psi_T \coloneqq \prod_{s=0}^{T-1} \{1 - \varphi_s(\mathbf{H}_s, \mathbf{P}_{s+1})\}$. In the payoff formulas below, for $t < T$, δ_t is evaluated at $(\mathbf{H}_t, \mathbf{P}_{t+1})$, while δ_T is evaluated at $(\mathbf{H}_T, \mathbf{P}_T)$.

Omitting the arguments of the strategy functions, the buyer's and the seller's ex ante expected payoffs are

$$U_B\{(\varphi, \delta), p\} \coloneqq \mathrm{E}\left[\sum_{t=0}^{T} \psi_t \left\{-L(\theta, \delta_t) - \sum_{s=1}^{t} P_s\right\}\right]$$

and

$$U_S\left\{(\varphi,\delta),p\right\} \coloneqq \mathrm{E}\left\{\sum_{t=1}^{T}\psi_t\left(\sum_{s=1}^{t}P_s-\sum_{s=1}^{t}c_s\right)\right\}.$$

Proof of Proposition 1. Fix φ and δ. By Assumption 1, for every t and $\mathbf{h}_t$,

$$\mathrm{E}\left[L\{\theta,\delta_t^0(\mathbf{h}_t)\}+\sum_{s=1}^{t}c_s \mid \mathbf{H}_t=\mathbf{h}_t\right] \leq \mathrm{E}\left[L\{\theta,\delta_t(\mathbf{h}_t)\}+\sum_{s=1}^{t}c_s \mid \mathbf{H}_t=\mathbf{h}_t\right].$$

Multiplying both sides by $\psi_t(\mathbf{h}_t)\geq 0$ and integrating over the distribution of $\mathbf{H}_t$, using iterated expectations, yields

$$\mathrm{E}\left(\psi_t(\mathbf{H}_t)\left[L\{\theta,\delta_t^0(\mathbf{H}_t)\}+\sum_{s=1}^{t}c_s\right]\right) \leq \mathrm{E}\left(\psi_t(\mathbf{H}_t)\left[L\{\theta,\delta_t(\mathbf{H}_t)\}+\sum_{s=1}^{t}c_s\right]\right).$$

Summing over t gives $r(\varphi,\delta^0)\leq r(\varphi,\delta)$. Therefore δ^0 is optimal.

Proof of Proposition 2. Fix δ^0 from Proposition 1. For each t and $\mathbf{h}_t$, let $W_t(\mathbf{h}_t)$ denote the minimal conditional expected social loss from period t onward, given $\mathbf{H}_t=\mathbf{h}_t$, when terminal decisions are taken according to δ^0.

At $t=T$, stopping is forced, so $W_T(\mathbf{h}_T)=\rho_T(\mathbf{h}_T)+\sum_{s=1}^{T}c_s=V_T(\mathbf{h}_T)$.

Now fix $t<T$. Conditional on $\mathbf{H}_t=\mathbf{h}_t$, the planner has two options. If the planner stops immediately, the induced loss is $\rho_t(\mathbf{h}_t)+\sum_{s=1}^{t}c_s$. If the planner continues, then after observing X_{t+1} the planner can attain $W_{t+1}(\mathbf{H}_{t+1})$, so the induced conditional expected loss is $\mathrm{E}\{W_{t+1}(\mathbf{H}_{t+1})\mid \mathbf{H}_t=\mathbf{h}_t\}$. Hence

$$W_t(\mathbf{h}_t)=\min\left[\rho_t(\mathbf{h}_t)+\sum_{s=1}^{t}c_s,\ \mathrm{E}\left\{W_{t+1}(\mathbf{H}_{t+1})\mid \mathbf{H}_t=\mathbf{h}_t\right\}\right].$$

By backward induction, $W_t(\mathbf{h}_t)=V_t(\mathbf{h}_t)$ for every t and $\mathbf{h}_t$. Therefore any stopping rule φ^0 satisfying the comparison in the statement attains V_t at each history, and $r(\varphi^0,\delta^0)=V_0(\mathbf{h}_0)\leq r(\varphi,\delta^0)$ for every stopping rule φ. Together with Proposition 1, (φ^0,δ^0) solves the social planner's problem.

Proof of Theorem 1. Step 1 (buyer optimality). The buyer's ex ante payoff from following (φ^*,δ^*) against p^* equals $-\rho_0$: $U_B\left\{(\varphi^*,\delta^*),p^*\right\}=-\rho_0$. To see this, if the buyer stops immediately, he plays a Bayes rule under the prior and obtains payoff $-\rho_0$. Otherwise the buyer purchases the first signal and then follows the efficient continuation policy, so his expected loss-plus-sampling-cost equals V_0, and paying the access fee $S=\rho_0-V_0$ yields payoff $-V_0-S=-V_0-(\rho_0-V_0)=-\rho_0$.

Fix p^* and let (φ, δ) be any buyer strategy. The strategy induces a stopping time τ taking values in $\{0, 1, \ldots, T\}$. Under p^*, the buyer's ex ante payoff is $U_B\{(\varphi, \delta), p^*\} = -\mathrm{E}[L(\theta, \delta_\tau) + \sum_{s=1}^{\tau} c_s + S1\{\tau \geq 1\}]$. Write $A \coloneqq \{\tau = 0\}$ and $A^c \coloneqq \{\tau \geq 1\}$. On A, the buyer stops at $t = 0$ after observing p_1 and (possibly) using private randomization, but neither the seller's posted price p_1 nor the buyer's $t = 0$ randomization depends on θ (there is no private information at $t = 0$). Hence conditioning on A does not change the distribution of θ. Let $\sigma_A \in \Delta(\mathcal{A})$ be the distribution over terminal actions induced by δ_0 conditional on A. Then, when $\Pr(A) > 0$,

$$\begin{aligned}\mathrm{E}\{L(\theta, \delta_0) \mid A\} &= \mathrm{E}\{L(\theta, \sigma_A) \mid A\} = \mathrm{E}\{L(\theta, \sigma_A)\} \\ &\geq \min_{\sigma \in \Delta(\mathcal{A})} \mathrm{E}\{L(\theta, \sigma)\} = \rho_0.\end{aligned}$$

On A^c, the buyer purchases the first signal before observing its realization. The event A^c is measurable with respect to date-0 information (the posted price p_1^* and the buyer's private randomization) and hence is independent of $\mathbf{H}_1 = X_1$. Moreover, conditional on $(A^c, \mathbf{H}_1)$, any continuation behavior from period 1 onward yields expected loss-plus-sampling-cost at least $V_1(\mathbf{H}_1)$ by the definition of V_1, so, when $\Pr(A^c) > 0$, $\mathrm{E}\{L(\theta, \delta_\tau) + \sum_{s=1}^{\tau} c_s \mid \mathbf{H}_1, A^c\} \geq V_1(\mathbf{H}_1)$ a.s. on A^c. Taking conditional expectations and using the tower property yields

$$\begin{aligned}\mathrm{E}\left\{L(\theta, \delta_\tau) + \sum_{s=1}^{\tau} c_s \mid A^c\right\} &= \mathrm{E}\left[\mathrm{E}\left\{L(\theta, \delta_\tau) + \sum_{s=1}^{\tau} c_s \mid \mathbf{H}_1, A^c\right\} \middle| A^c\right] \\ &\geq \mathrm{E}\{V_1(\mathbf{H}_1) \mid A^c\} = \mathrm{E}\{V_1(\mathbf{H}_1)\},\end{aligned}$$

where the last equality uses $\mathbf{H}_1 \perp A^c$. Combining these inequalities gives

$$U_B\{(\varphi, \delta), p^*\} \leq -\Pr(A)\rho_0 - \Pr(A^c)\left[\mathrm{E}\{V_1(\mathbf{H}_1)\} + S\right].$$

Plugging in the definition $S = \rho_0 - V_0$ with $V_0 = \min[\rho_0, \mathrm{E}\{V_1(\mathbf{H}_1)\}]$ gives

$$U_B\{(\varphi, \delta), p^*\} \leq -\rho_0 - \Pr(A^c)\left(\mathrm{E}\{V_1(\mathbf{H}_1)\} - \min[\rho_0, \mathrm{E}\{V_1(\mathbf{H}_1)\}]\right) \leq -\rho_0,$$

since $\mathrm{E}\{V_1(\mathbf{H}_1)\} - \min[\rho_0, \mathrm{E}\{V_1(\mathbf{H}_1)\}] \geq 0$. Therefore no deviation yields a higher payoff, and (φ^*, δ^*) is a best response to p^*.

Step 2 (seller optimality). Fix (φ^*, δ^*) and let $\tilde{p} = (\tilde{p}_t)_{t=1}^T$ be any seller pricing strategy. Let $N \in \{0, 1, \ldots, T\}$ denote the number of signals actually purchased when the buyer plays (φ^*, δ^*) against $\tilde{p}$.

After any history $(\mathbf{h}_t, \mathbf{p}_{t+1})$ with $\mathbf{p}_{t+1} \neq (p_1^*, \ldots, p_{t+1}^*)$, the buyer stops immediately and

does not purchase signal $t+1$. Hence, on any realized history with $N = k \geq 1$, the prices of the purchased signals must be exactly $p_1^*, p_2^*, \ldots, p_k^*$. Therefore, on any such realized history, the seller's payoff is $\sum_{t=1}^{k} p_t^* - \sum_{t=1}^{k} c_t = \left\{(c_1 + S) + \sum_{t=2}^{k} c_t\right\} - \sum_{t=1}^{k} c_t = S$. If instead $N = 0$, then the seller's payoff is 0.

Thus, for every realized history, the seller's payoff under $\tilde{p}$ is either 0 or S, and hence is at most S, since $S = \rho_0 - V_0 \geq 0$ because $V_0 \leq \rho_0$ by definition. It follows that $U_S\{(\varphi^*, \delta^*), \tilde{p}\} \leq S$. Under p^*, the seller's equilibrium payoff is S. Therefore p^* is a best response to (φ^*, δ^*).

Since each player is best responding, $\{(\varphi^*, \delta^*), p^*\}$ is a Nash equilibrium.

Proof of Theorem 2. Fix $t \in \{0, 1, \ldots, T\}$, $\mathbf{h}_t$, and $\mathbf{p}_t$. For the candidate strategy profile, all continuation strategies from this point onward depend on $\mathbf{h}_t$ but not on $\mathbf{p}_t$; hence, in the constructive part of the proof, I suppress $\mathbf{p}_t$ from the notation for continuation values.

Define the shifted planner value function $\widetilde{V}_t(\mathbf{h}_t) \coloneqq V_t(\mathbf{h}_t) - \sum_{s=1}^{t} c_s$. By the definition of V_t preceding Proposition 2, we have $\widetilde{V}_T(\mathbf{h}_T) = \rho_T(\mathbf{h}_T)$ and, for $t < T$,

$$\widetilde{V}_t(\mathbf{h}_t) = \min\left[\rho_t(\mathbf{h}_t),\ c_{t+1} + \mathrm{E}\left\{\widetilde{V}_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}\right]. \tag{B1}$$

Moreover, by Definition 1,

$$S_t(\mathbf{h}_t) = \rho_t(\mathbf{h}_t) + \sum_{s=1}^{t} c_s - V_t(\mathbf{h}_t) = \rho_t(\mathbf{h}_t) - \widetilde{V}_t(\mathbf{h}_t). \tag{B2}$$

For $t < T$, let $B_t(\mathbf{h}_t)$ and $\Pi_t(\mathbf{h}_t)$ denote, respectively, the buyer's and seller's continuation payoffs under the candidate profile in the subgame that starts at the seller's period-$t+1$ pricing node after history $\mathbf{H}_t = \mathbf{h}_t$, both net of past payments and costs. For $t = T$, let $B_T(\mathbf{h}_T)$ and $\Pi_T(\mathbf{h}_T)$ denote the same objects in the subgame that starts at the buyer's terminal decision node after $\mathbf{H}_T = \mathbf{h}_T$.

I prove by backward induction that, for every t and $\mathbf{h}_t$, $B_t(\mathbf{h}_t) = -\rho_t(\mathbf{h}_t)$ and $\Pi_t(\mathbf{h}_t) = S_t(\mathbf{h}_t)$, and that the seller pricing rule in the statement is optimal in every continuation game.

Step 1: terminal period. At $t = T$, no further trade is possible. The buyer stops and chooses a Bayes act, so $B_T(\mathbf{h}_T) = -\rho_T(\mathbf{h}_T)$. The seller's continuation payoff is $\Pi_T(\mathbf{h}_T) = 0$. Since $\widetilde{V}_T(\mathbf{h}_T) = \rho_T(\mathbf{h}_T)$, (B2) implies $S_T(\mathbf{h}_T) = 0$, hence $\Pi_T(\mathbf{h}_T) = S_T(\mathbf{h}_T)$.

Step 2: induction step. Fix $t < T$ and suppose that for all $\mathbf{h}_{t+1}$, $B_{t+1}(\mathbf{h}_{t+1}) = -\rho_{t+1}(\mathbf{h}_{t+1})$ and $\Pi_{t+1}(\mathbf{h}_{t+1}) = S_{t+1}(\mathbf{h}_{t+1})$.

Consider first the buyer's decision at $(\mathbf{h}_t, p_{t+1})$. If the buyer stops, his continuation payoff is $-\rho_t(\mathbf{h}_t)$. If he buys, then by the induction hypothesis his continuation payoff is

$-p_{t+1} - \mathrm{E}\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}$. Therefore

$$\widehat{B}_t(\mathbf{h}_t, p_{t+1}) = \max\left[-\rho_t(\mathbf{h}_t),\ -p_{t+1} - \mathrm{E}\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}\right],$$

and the buyer purchases if and only if

$$p_{t+1} \leq \bar{p}_{t+1}(\mathbf{h}_t) \coloneqq \rho_t(\mathbf{h}_t) - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}.$$

Now consider the seller's pricing problem at history $\mathbf{h}_t$. If she wants trade, then any price below $\bar{p}_{t+1}(\mathbf{h}_t)$ is strictly dominated by $\bar{p}_{t+1}(\mathbf{h}_t)$, while any price above $\bar{p}_{t+1}(\mathbf{h}_t)$ induces stopping. Hence the maximal payoff from inducing trade is attained at $p_{t+1} = \bar{p}_{t+1}(\mathbf{h}_t)$, which yields

$$\begin{aligned}
&\bar{p}_{t+1}(\mathbf{h}_t) - c_{t+1} + \mathrm{E}\left\{\Pi_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\} \\
&= \rho_t(\mathbf{h}_t) - \mathrm{E}\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} - c_{t+1} \\
&\quad + \mathrm{E}\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} \\
&= \rho_t(\mathbf{h}_t) - c_{t+1} - \mathrm{E}\left\{\widetilde{V}_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\},
\end{aligned} \tag{B3}$$

where the second equality uses (B2). If instead the seller induces stopping by choosing any $p_{t+1} > \bar{p}_{t+1}(\mathbf{h}_t)$, her continuation payoff is 0. Therefore

$$\Pi_t(\mathbf{h}_t) = \max\left[0,\ \rho_t(\mathbf{h}_t) - c_{t+1} - \mathrm{E}\left\{\widetilde{V}_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}\right]. \tag{B4}$$

Using (B1),

$$\widetilde{V}_t(\mathbf{h}_t) = \min\left[\rho_t(\mathbf{h}_t),\ c_{t+1} + \mathrm{E}\left\{\widetilde{V}_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}\right],$$

so (B4) is exactly $\Pi_t(\mathbf{h}_t) = \rho_t(\mathbf{h}_t) - \widetilde{V}_t(\mathbf{h}_t) = S_t(\mathbf{h}_t)$. Hence the seller's continuation payoff is $S_t(\mathbf{h}_t)$.

This also identifies the optimal pricing rule. If $S_t(\mathbf{h}_t) > 0$, then $\rho_t(\mathbf{h}_t) > c_{t+1} + \mathrm{E}\{\widetilde{V}_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}$, so the quantity in (B3) is strictly positive and the seller optimally induces trade by posting the maximal acceptable price

$$p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t) = \rho_t(\mathbf{h}_t) - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}.$$

If $S_t(\mathbf{h}_t) = 0$, then the seller's maximal continuation payoff is 0. In particular, the price

$$p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t, \mathbf{p}_t) = \rho_t(\mathbf{h}_t) - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\} + 1$$

is optimal and induces stopping. More generally, any measurable price strictly above the buyer's threshold is also optimal on such histories.

At the prescribed equilibrium price, the buyer's continuation payoff equals $-\rho_t(\mathbf{h}_t)$ in both cases: if $S_t(\mathbf{h}_t) > 0$, the seller posts the threshold price and the buyer is indifferent, so $B_t(\mathbf{h}_t) = \widehat{B}_t\left\{\mathbf{h}_t, p_{t+1}^{\mathrm{SPE}}(\mathbf{h}_t)\right\} = -\rho_t(\mathbf{h}_t)$; if $S_t(\mathbf{h}_t) = 0$, the buyer stops and again receives $-\rho_t(\mathbf{h}_t)$.

Step 3: subgame perfection, efficiency, and payoff uniqueness.

By the induction argument, after every seller-node history $(\mathbf{h}_t, \mathbf{p}_t)$ the seller's prescribed price is optimal given the buyer's continuation strategy, and after every buyer-node history $(\mathbf{h}_t, \mathbf{p}_{t+1})$ the buyer's prescribed action is optimal given the seller's continuation pricing rule. At terminal histories, $\delta^{\mathrm{SPE}} = \delta^0$ is Bayes optimal by Assumption 1. Hence the strategy profile is an SPE.

The induced stopping rule continues exactly on histories with $S_t(\mathbf{h}_t) > 0$ and stops on histories with $S_t(\mathbf{h}_t) = 0$. By (B1),

$$S_t(\mathbf{h}_t) > 0 \iff \rho_t(\mathbf{h}_t) + \sum_{s=1}^{t} c_s > \mathrm{E}\left\{V_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}.$$

Thus the equilibrium stopping rule coincides with the efficient stopping rule on all strict-comparison histories and selects stopping on equality histories, which is also efficient. Since $\delta^{\mathrm{SPE}} = \delta^0$, terminal actions are efficient as well. Therefore the equilibrium is efficient.

Finally, the ex ante continuation payoffs in the constructed equilibrium are $\Pi_0(\mathbf{h}_0) = S_0(\mathbf{h}_0) = S$ and $B_0(\mathbf{h}_0) = -\rho_0$.

To prove uniqueness of the SPE continuation payoff vector before each quote, consider any SPE. For each seller-node history $(\mathbf{h}_t, \mathbf{p}_t)$, let $B_t^*(\mathbf{h}_t, \mathbf{p}_t)$ and $\Pi_t^*(\mathbf{h}_t, \mathbf{p}_t)$ denote the buyer's and seller's continuation payoffs, and for each buyer-node history $(\mathbf{h}_t, \mathbf{p}_{t+1})$, let $\widehat{B}_t^*(\mathbf{h}_t, \mathbf{p}_{t+1})$ denote the buyer's continuation payoff after the current quote. I claim that for every seller-node history, $B_t^*(\mathbf{h}_t, \mathbf{p}_t) = -\rho_t(\mathbf{h}_t)$ and $\Pi_t^*(\mathbf{h}_t, \mathbf{p}_t) = S_t(\mathbf{h}_t)$, and for every buyer-node history with $t < T$,

$$\widehat{B}_t^*(\mathbf{h}_t, \mathbf{p}_{t+1}) = \max\left[-\rho_t(\mathbf{h}_t),\ -p_{t+1} - \mathrm{E}\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\}\right].$$

The base case $t = T$ is immediate. For the induction step, use the induction hypothesis at $t+1$. At a buyer node after quote p_{t+1}, stopping yields $-\rho_t(\mathbf{h}_t)$, while buying yields

$$\begin{aligned} &-p_{t+1} + \mathrm{E}\left\{B_{t+1}^*(\mathbf{H}_{t+1}, \mathbf{p}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\} \\ &= -p_{t+1} - \mathrm{E}\left\{\rho_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\right\}. \end{aligned}$$

So the buyer's acceptance threshold is $\bar{p}_{t+1}(\mathbf{h}_t)$. I now distinguish two cases.

If $S_t(\mathbf{h}_t) > 0$, then

$$\bar{p}_{t+1}(\mathbf{h}_t) - c_{t+1} + \mathrm{E}\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} = S_t(\mathbf{h}_t) > 0.$$

Any quote $p_{t+1} > \bar{p}_{t+1}(\mathbf{h}_t)$ yields seller payoff 0, so it cannot be optimal. Any accepted quote $p_{t+1} < \bar{p}_{t+1}(\mathbf{h}_t)$ cannot be optimal either, since for sufficiently small $\varepsilon > 0$, the higher quote $p_{t+1} + \varepsilon < \bar{p}_{t+1}(\mathbf{h}_t)$ is still strictly accepted and yields strictly higher seller payoff. Finally, if the seller quotes $p_{t+1} = \bar{p}_{t+1}(\mathbf{h}_t)$ and the buyer buys with probability $q < 1$, the seller's payoff is $qS_t(\mathbf{h}_t)$, and she can deviate to $p'_{t+1} = \bar{p}_{t+1}(\mathbf{h}_t) - \varepsilon$ for some $0 < \varepsilon < (1-q)S_t(\mathbf{h}_t)$; the buyer then strictly prefers to buy, and the seller obtains payoff $S_t(\mathbf{h}_t) - \varepsilon > qS_t(\mathbf{h}_t)$. Therefore every SPE must have trade with probability one at the threshold price $p_{t+1} = \bar{p}_{t+1}(\mathbf{h}_t)$. Hence

$$\Pi^*_t(\mathbf{h}_t, \mathbf{p}_t) = \bar{p}_{t+1}(\mathbf{h}_t) - c_{t+1} + \mathrm{E}\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} = S_t(\mathbf{h}_t),$$

and $B^*_t(\mathbf{h}_t, \mathbf{p}_t) = -\rho_t(\mathbf{h}_t)$.

If instead $S_t(\mathbf{h}_t) = 0$, then

$$\bar{p}_{t+1}(\mathbf{h}_t) - c_{t+1} + \mathrm{E}\{S_{t+1}(\mathbf{H}_{t+1}) \mid \mathbf{H}_t = \mathbf{h}_t\} \leq 0.$$

Any accepted quote $p_{t+1} < \bar{p}_{t+1}(\mathbf{h}_t)$ yields strictly negative seller payoff and hence cannot be optimal. Any quote $p_{t+1} > \bar{p}_{t+1}(\mathbf{h}_t)$ induces stopping and yields seller payoff 0. At the threshold price $p_{t+1} = \bar{p}_{t+1}(\mathbf{h}_t)$, trade gives the seller a nonpositive payoff and stopping gives 0, and in either case the buyer's continuation payoff is $-\rho_t(\mathbf{h}_t)$. Thus $\Pi^*_t(\mathbf{h}_t, \mathbf{p}_t) = 0 = S_t(\mathbf{h}_t)$ and $B^*_t(\mathbf{h}_t, \mathbf{p}_t) = -\rho_t(\mathbf{h}_t)$.

This proves uniqueness of the SPE continuation payoff vector before each quote, and of the buyer's continuation payoff after every quote.

References

Admati, Anat R., and Paul Pfleiderer. 1986. "A Monopolistic Market for Information." *Journal of Economic Theory*, 39(2): 400–438. https://doi.org/10.1016/0022-0531(86)90052-9.

Admati, Anat R., and Paul Pfleiderer. 1990. "Direct and Indirect Sale of Information." *Econometrica*, 58(4): 901–928. https://doi.org/10.2307/2938355.

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review*, 108(1): 1–48. https://doi.org/10.1257/aer.20161079.

Bergemann, Dirk, and Alessandro Bonatti. 2019. "Markets for Information: An Introduction." *Annual Review of Economics*, 11(1): 85–107. https://doi.org/10.1146/annurev-economics-080315-015439.

Birks, Ian, and Larisa Nagra Singh. 2011. "FSPs: Ranging from Tactical to Strategic Models." White paper, Quintiles. https://www.iqvia.com/-/media/library/white-papers/fsps-ranging-from-tactical-to-strategic-models.pdf.

Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. Jerzy Neyman, 93–102. Berkeley: University of California Press. https://doi.org/10.1525/9780520411586-009.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics*, 24(2): 265–272. https://doi.org/10.1214/aoms/1177729032.

Bonatti, Alessandro, Munther Dahleh, Thibaut Horel, and Amir Nouripour. 2024. "Selling Information in Competitive Environments." *Journal of Economic Theory*, 216: 105779. https://doi.org/10.1016/j.jet.2023.105779.

Chen, Chia-Hui, Junichiro Ishida, and Wing Suen. 2026. "Dynamic Pricing of Information: Belief Divergence and Surplus Extraction." Institute of Social and Economic Research, The University of Osaka Discussion Paper 1294. Revised June 11, 2026. https://drive.google.com/file/d/1TalcMYo-sZtAGEhv8yFpQVYp7CPLF1ZA/view.

Eső, Péter, and Balázs Szentes. 2007. "The Price of Advice." *RAND Journal of Economics*, 38(4): 863–880. https://doi.org/10.1111/j.0741-6261.2007.00116.x.

FDA. 2018. "Step 3: Clinical Research." https://www.fda.gov/patients/drug-development-process/step-3-clinical-research.

Hörner, Johannes, and Andrzej Skrzypacz. 2016. "Selling Information." *Journal of Political Economy*, 124(6): 1515–1562. https://doi.org/10.1086/688874.

ICON plc. 2025. "Form 20-F (Annual Report) for the Fiscal Year Ended December 31, 2024." https://www.sec.gov/Archives/edgar/data/1060955/000106095525000016/iclr-20241231.htm.

Kruse, Thomas, and Philipp Strack. 2015. "Optimal Stopping with Private Information." *Journal of Economic Theory*, 159: 702–727. https://doi.org/10.1016/j.jet.2015.03.001.

Lehrer, Ehud, and Tao Wang. 2024. "The Value of Information in Stopping Problems." *Economic Theory*, 78(2): 619–648. https://doi.org/10.1007/s00199-023-01543-8.

Liu, Shuze, Weiran Shen, and Haifeng Xu. 2021. "Optimal Pricing of Information." arXiv Preprint 2102.13289v3. https://arxiv.org/abs/2102.13289v3.

Mekonnen, Teddy, Zeky Murra-Anton, and Bobak Pakzad-Hurson. 2025. "Persuaded Search." *Journal of Political Economy*, 133(10): 3167–3207. https://doi.org/10.1086/736765.

Morris, Stephen, and Philipp Strack. 2019. "The Wald Problem and the Relation of Sequential Sampling and Ex-Ante Information Costs." Working paper, February 18. https://doi.org/10.2139/ssrn.2991567.

Oi, Walter Y. 1971. "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly." *Quarterly Journal of Economics*, 85(1): 77–96. https://doi.org/10.2307/1881841.

PPD. 2026. "Meet Your Timelines with PPD Functional Service Partnership (FSP) Solutions." Brochure, Thermo Fisher Scientific. https://www.ppd.com/wp-content/uploads/2023/07/PPD-FSP-Solutions-Overview.pdf.

Rivera Mora, Ernesto, and Nicolás Rodríguez. 2026. "Information Selling under Prior Disagreement." Working paper, April. https://doi.org/10.2139/ssrn.6945218.

Rodríguez Olivera, Rosina. 2024. "Strategic Incentives and the Optimal Sale of Information." *American Economic Journal: Microeconomics*, 16(2): 296–353. https://doi.org/10.1257/mic.20210372.

Wald, Abraham. 1945. "Sequential Tests of Statistical Hypotheses." *Annals of Mathematical Statistics*, 16(2): 117–186. https://doi.org/10.1214/aoms/1177731118.