

Sequentially Optimal Pricing under Worst-Case Information

ZIHAO LI* JONATHAN LIBGOBER** XIAOSHENG MU***

*Columbia University

**University of Southern California

***Princeton University

September 2, 2026

ABSTRACT. A seller sells an object over time but is uncertain how the buyer learns her willingness-to-pay. Each period, the seller chooses a price to maximize expected continuation profit against *worst-case information arrival*, doing so under *limited commitment*. We formulate the worst case *sequentially* and characterize an essentially unique equilibrium. We further show that, under mild conditions on the prior distribution, the equilibrium profit coincides with the profit guaranteed by the equilibrium price path at any on-path history, even when the worst-case is not considered sequentially.

KEYWORDS. Worst-case information, Limited commitment, Safe solution.

CONTACT. ZL3366@COLUMBIA.EDU, LIBGOBER@USC.EDU, XMU@PRINCETON.EDU. A longer version of this paper, available at www.jonlib.com, includes additional results on alternative formulations of the worst-case objective. We thank Dirk Bergemann, Alex Bloedel, Daniel Clark, Francesco Fabbri, Drew Fudenberg, Johannes Hörner, Navin Kartik, Deniz Kattwinkel, Peter Klibanoff, George Mailath, Rob Metcalfe, Mathias Nunez, Alessandro Pavan, João Ramos, Andy Skrzypacz, Balázs Szentes and Kai Hao Yang for helpful comments and seminar audiences at Boston University, Brown, Harvard/MIT, LSE, NASMES, Rochester, SITE: Dynamic Games, Stony Brook, UChicago, UPenn, USC and Yale for feedback. The second author thanks the hospitality of the Cowles Foundation and Yale University, which hosted him during part of this research.

1. INTRODUCTION

This paper studies the following situation. A seller aims to sell one unit of an indivisible good to a buyer, but faces *limited commitment*, only able to commit to a price within the period it is offered. The buyer, in turn, may not initially know her willingness-to-pay, but can learn about it through some (potentially complex) process while considering purchase. A defining feature of our exercise is that we consider a seller lacking the knowledge necessary to form a prior over how information arrives. Prices are therefore chosen to maximize expected profit against worst-case information arrival. We are particularly interested in situations where the seller is forward-looking, thus reoptimizing the price each period to maximize payoffs *from that time onward*.

The possibility of information arrival has clear economic significance, as buyers may lack precise knowledge of their willingness-to-pay for many reasons. Consider, for example, a new parent negotiating to buy a house from a previous homeowner who has already vacated.¹ Neighborhood characteristics—such as average annual NO_x levels, risks from natural disasters, or school quality—may not be immediately apparent.² A first-time buyer may not even initially understand how such characteristics *should* translate into a true willingness-to-pay. Accordingly, how a particular buyer learns may be highly idiosyncratic, and could vary dramatically in hard-to-predict ways. The buyer might be perfectly informed at the outset—or not. She might consult family or various online resources, or even learn about when further information will arrive.

In principle, the possibility of information arrival could fundamentally alter behavior in applications where the seller lacks commitment to future prices. This claim follows from the standard intuition that the degree to which delay selects low willingness-to-pay buyers shapes equilibrium pricing (see Section 1.3 for discussion). While incentives to delay under a given price path are determined by expected surplus absent information, this logic can break down when information arrival is possible. A buyer who expects to value the product highly might nevertheless delay to become more certain that this is the case. Past work has convincingly demonstrated the significance of limited commitment, but relatively little work has considered the implications of rich possibilities for information arrival.³

¹Such negotiations often lack mechanisms to avoid price revisions. The sale of land is the leading example in Coase (1972). Han and Strange (2015) discuss recent evidence on post-match price revisions in housing markets.

²Information availability about idiosyncratic value-relevant characteristics is particularly well-documented in housing markets. Fairweather et al. (2023) conduct an experiment showing that providing information about flood risk significantly changes buyer behavior. Ainsworth et al. (2023) show that households often lack accurate information about local school quality, and that belief accuracy predicts child achievement. Bergman et al. (2020) show experimentally that providing school information influences household location choices.

³As we discuss in Section 1.3, existing precedents impose particular structure on information arrival. As Pavan (2017) notes: “The literature on limited commitment has made important progress in recent years.... However, this literature assumes information is static, thus abstracting from the questions at the heart of the dynamic mechanism design literature. I expect interesting new developments to come out from combining the two literatures.” We view our

And indeed, these implications could be quite dramatic. Previewing an observation to come, in this paper we show that constant price paths can be sustained and multiple equilibria can arise if arbitrary *exogenous* and commonly-known information arrival processes are possible—a stark departure from outcomes when willingness-to-pay is known *ex ante* (see Section 5.1). We view this as a proof of concept that some structure is necessary to derive interpretable pricing heuristics or to characterize equilibrium properties. Yet in many cases such restrictions lack economic justification, motivating a different perspective to serve as a useful baseline.

Our goal is to obtain sharp equilibrium predictions under a worst-case objective when the seller lacks a prior over how the buyer learns. Typically, worst-case objectives are motivated economically by a lack of confidence in forming a prior over some aspect of the relevant environment—in our case, how the buyer learns. Such concerns also appear relevant under limited commitment, but the lack of commitment introduces conceptual difficulties in formulating a worst-case objective.

These difficulties arise because limited commitment requires the seller to consider the worst-case each time a price is chosen, rather than once. Past work on single-agent decision problems has shown that this requirement introduces the potential for inconsistencies in the worst-case solutions. To illustrate, suppose that at time 2 the seller fears that the buyer will perfectly learn her value at time 10, encouraging a “wait-and-see” strategy. Yet once time 10 arrives, such information might not be the worst for the seller; keeping a buyer with a value slightly above the price ignorant could deter purchase. The worst-case objective at time 2 would thus conflict with that at time 10. Dynamic inconsistencies in maxmin models such as this have been extensively studied in decision theory, where there is some debate about whether dynamic consistency is *per se* desirable.⁴

Setting this debate aside, our setting is fundamentally different from single-agent dynamic worst-case problems (such as Auster et al. (2022); Malladi (2023)). Relative to the single-agent case, the presence of a buyer introduces new conceptual issues for the seller’s objective.

Most of these issues stem from the requirement that buyer behavior be sequentially rational given the information arrival process, rather than being selected directly to minimize seller profit. Without rationality constraints, the worst-case buyer strategy would be trivial: never purchase. Imposing rationality constraints on buyer behavior is also standard in other work employing worst-case objectives *under commitment*. There, however, the conceptual difficulties this requirement introduces are more limited, because the buyer (and, say, an adversarial “Nature” that makes uncertainty resolve toward the worst case) can take the seller’s choices as fixed.

contribution in the spirit of this agenda.

⁴For example, Al-Najjar and Weinstein (2009) document behavioral anomalies that can emerge in dynamic maxmin models without dynamic consistency, while Siniscalchi (2009) argues that several of these anomalies may be natural. See also Epstein and Schneider (2003) for discussion of why dynamic consistency may be desirable. Related inconsistencies can arise in sequential games with ambiguity even outside of worst-case information; see Battigalli et al. (2019a,b) for versions under *smooth-ambiguity preferences*, which can approximate maxmin preferences.

Once commitment is relaxed, however, buyer assumptions about *how* prices are determined fundamentally influence outcomes (as do, of course, the seller’s assumptions about the buyer). Simply stating how uncertainty resolves is not enough to make equilibrium outcomes well-defined without also considering the buyer’s problem.

This makes dynamic inconsistency more consequential in the strategic setting, since inconsistencies in worst-case assessments over time can interact with the assumptions used to determine the buyer’s strategy. The sale of a durable good provides a canonical environment in which to study these issues, since the seller chooses a one-dimensional price while the buyer faces a stopping problem. This setting is the natural benchmark because it isolates the key strategic forces that make the consistency problem substantive while remaining sufficiently structured to admit a sharp equilibrium characterization. We use it to obtain sharp predictions despite general information arrival.

1.1. Our Approach

In the simple interaction we study, the seller offers the object at a price chosen each period, and the buyer—after receiving her current information—either accepts or rejects the offer. The interaction continues until a terminal date or indefinitely. This environment reduces to a textbook bargaining model if the buyer knows her willingness-to-pay from the outset; under a Bayesian approach, the seller would form a prior over the buyer’s information arrival process.

To formulate the worst-case coherently, we make the following pair of assumptions:

1. Whenever the buyer decides whether to purchase, her information set includes only *realized information*—the sequence of realized information structures and signals—generated by the information arrival process up to that history, not the full process or its future realizations.
2. The seller posits that the buyer’s information arrival process minimizes the seller’s continuation profit *at every history*.

In equilibrium, buyer and seller strategies are optimal given the resulting process.

We call such an information arrival process *sequentially worst-case*. The first assumption ensures that continuation play at different histories is determined by one dynamically coherent information-arrival process together with a fixed buyer strategy, rather than by potentially disconnected history-by-history worst cases. The second requires that the process is worst case each time a price is chosen, rather than only *ex ante*. This feature captures limited commitment because the seller re-evaluates the worst case each time he re-optimizes the price.

This approach contrasts with common formulations in robust mechanism design, which consider designers who maximize a payoff guarantee over a set of possible environments. While

our seller cannot freely optimize over payoff guarantees—since prices are restricted to form an equilibrium—our formulation *does* admit a payoff guarantee interpretation via the game’s normal-form representation. There, the seller chooses a complete pricing strategy at date 0, while “Nature” chooses a complete history-dependent information arrival process. Future prices and information realizations are revealed to the buyer only as play unfolds.⁵ Sequentially worst-case information arises from the Nash equilibrium of this zero-sum game between the seller and Nature.

1.2. Our Results

Our approach yields a dynamically coherent worst-case information-arrival process where the continuation process that is worst case at any history coincides with the process conjectured at earlier histories. Our two main results are the following:

1. **Sharp characterization of equilibrium pricing.** The equilibrium is essentially unique,⁶ with the following property: prices decline over time, and in any period the buyer’s decision is identical to the one she would make if she considered only the information available in that period, as though no further information would arrive. The sharpness of this result contrasts starkly to the case where particular, commonly-known processes are assumed.
2. **Moving beyond sequentiality using safe solutions.** Our benchmark restricts the buyer to conditioning on realized information. We ask whether the equilibrium price path achieves the same payoff guarantee if the buyer instead knew the entire information-arrival process, including how she will learn in the future. For a broad class of value distributions, it does. We call such a price path a *safe solution*.

Our equilibrium characterization shows that sequentially worst-case information arrival minimizes the probability of sale *within each period* given the price path. The key intuition is that worst-case continuation information eliminates the option value of future learning when the buyer is indifferent between purchase and delay. As a result, the indifference point depends only on the conjectured price path. This reduction allows us to characterize equilibrium explicitly and obtain a structure of equilibrium price paths analogous to that which arises absent learning. In particular, the seller does not randomize on the equilibrium path, except possibly in the initial period.

With these observations in hand, the price path itself can be determined by appealing to backward inductive reasoning, an intuitive consequence of dynamic consistency. In the final period, worst-case information has the property that any buyer who does not purchase is indifferent

⁵If the buyer knows willingness-to-pay ex ante, the “secret price path” formulation reduces to the aforementioned textbook bargaining models. The normal form we outline thus generalizes this formulation.

⁶As in Fudenberg et al. (1985): actions are deterministic after possible seller randomization in the first period.

between purchasing and not, making her payoff the same regardless of her action. In particular, her payoff would be unchanged were she to instead *always* purchase, despite her breaking indifference by not purchasing. Hence, in the second-to-last period, the buyer behaves as if no future information will arrive. The worst-case information in that period thus simply minimizes the probability of purchase, and the same logic propagates backward.

In our model, this implies that a seller who evaluates the worst case at every history, and anticipates that the buyer makes decisions using only information available up to the current period, behaves as if the sale probability is minimized in each period given the pricing strategy. In this sense the worst-case is *myopic*. The reduction to a myopic worst-case implies that equilibrium pricing takes a simple form despite the potential complexity of information-arrival.

Our second main result shows that, for a large class of environments, sequentially worst-case information arrival minimizes the seller's profit at *all* on-path histories, even when the buyer's strategy is allowed to depend on the entire process, *fixing* the pricing strategy to be the one that arises in equilibrium. This result provides an additional layer of robustness: the seller's worst-case payoff is unchanged even when the buyer's strategy is allowed to condition on more than the realized history. This preserves a central appeal of the worst-case objective—the equilibrium strategy continues to perform well even when the informational possibilities under consideration are enlarged. Put differently, even if the seller were misspecified about the buyer's knowledge of the process itself, this misspecification could not reduce his payoff. We call any equilibrium pricing strategy with this property a *safe solution*.

Intuitively, providing more information later involves a tradeoff between (i) inducing more delay by higher-value buyers in earlier periods, and (ii) inducing more—or at least earlier—sales by lower-value buyers in later periods. The condition on the prior value distribution that we identify ensuring the price path is a safe solution—which we call *threshold-ratio monotonicity*—essentially requires that the increase in willingness-to-delay from future information is small relative to the additional sales such information generates. The condition includes, among others, uniform distributions and the power family (i.e. $F(v) = v^a$ for $a > 0$), two standard benchmark families.

This result is conceptually subtle because it requires the worst case to be evaluated *holding fixed* the seller's pricing strategy. This differs from the alternative approach in which one considers the worst-case equilibrium in the buyer-seller game that can emerge given an *exogenous* commonly-known information arrival process. It is possible that some other pricing strategy could arise in equilibrium under such a process, under which the seller would be worse off.⁷

⁷More generally, the possibility of shifting continuation surplus toward a delaying buyer creates scope for commonly-known information arrival to reduce seller profit further; Appendix Section C.1 develops this intuition and provides an explicit two-period uniform example illustrating how this can be done.

1.3. Relevant Literature

A major point of departure of our paper from the literature on informationally robust mechanism design⁸ is our focus on the equilibrium outcomes of a particular bargaining protocol (rather than allowing general mechanisms). To our knowledge, the two papers closest to this exercise are Libgober and Mu (2021) and Koh and Sanguanmoo (2025), as both consider stopping problems with worst-case information. Libgober and Mu (2021) study robust dynamic pricing *with commitment* so the seller faces a single ex-ante optimization problem; here, prices must instead remain optimal when future dates arrive, making repeated evaluation of the informational worst case itself an equilibrium object.⁹ In Koh and Sanguanmoo (2025), the agent’s stopping time determines the level of investment in a technology that can be either good or bad. Notably, this environment does *not* feature a difference between the commitment and non-commitment solutions—a stark contrast with durable goods sales. A defining feature for us is that the repeated evaluation of the worst case influences equilibrium.¹⁰

A less directly related literature considers mechanism design where agents—rather than the designer—have non-Bayesian preferences, including the maxmin case (Bose and Renou, 2014; Wolitzky, 2016; Di Tillio et al., 2017). The typical focus there is on how the *designer* can exploit such preferences; some papers explicitly examine exploiting dynamic inconsistency (Bose et al., 2006; Bose and Daripa, 2009).

Our model would reduce to the textbook durable goods monopoly without commitment (Fudenberg et al., 1985; Gul et al., 1986) if the buyer were to learn her realized willingness-to-pay in the initial period. While we use some technical results from this literature, our focus is largely orthogonal: the $\delta \rightarrow 1$ limit is not our main interest, unlike most papers in this area. A notable exception is Fuchs and Skrzypacz (2013), which characterizes *non-trivial* pricing dynamics with “frequent offers” and a finite time to trade.

Changes in preferences can alter the conclusions of the Coase conjecture literature (e.g., Ortner (2017, 2023); Acharya and Ortner (2017)), and information arrival can be interpreted as a form of preference change. Relatedly, Lomys (2018), Duraj (2020), and Laiho and Salmi (2020) study how Coasian dynamics are affected by buyer information, but under restrictive assumptions on type distributions or information arrival processes. These effects arise from the interaction

⁸For instance, Bergemann et al. (2017); Du (2018); Brooks and Du (2021, 2023); Deb and Roesler (2023). Other papers considering a worst-case objective without the focus on information include Carroll (2015, 2017); Lopomo et al. (2020); Che and Zhong (2022); He and Li (2022).

⁹The commitment case also differs economically: future prices are part of the seller’s initial optimization problem, rather than subject to optimality requirements when the future arrives. In this sense, the committed seller does not face competition with his future self.

¹⁰Other dynamic robust mechanism design papers include Chassang (2013); Penta (2015); Durandard et al. (2024), but in these and most others, the worst case is considered only once.

between information arrival and selection; the importance of the direction of selection for Coasian dynamics is highlighted in Tirole (2016). In contrast, one interpretation of our main result is that Coasian dynamics are fully restored against worst-case information. This restoration may fail with other worst-case formulations (see the longer working paper version of this paper).

Finally, our *safe solution* requirement connects to recent proposals to strengthen worst-case criteria—seeking not only optimality against a *single* worst case, but also good performance across a broader set of possibilities. Kambhampati (2025) pursues this goal, but evaluates performance against alternative possibilities from the *same* set as the one initially used to define worst-case performance. Ball and Kattwinkel (2025) also considers mechanism performance under ambiguity set expansions, as the safe solution requirement does, but focuses on local expansions—requiring near-optimal performance when Nature’s choices are “close” to those the designer entertains. While our definition of a safe solution focuses on the problem at hand to maintain focus, the same definition extends readily to any setting where robustness of a policy against a restricted class is sought, but then evaluated against an unrestricted class.

2. MODEL

We first present the basic primitives of the environment. We then turn to the mechanics of how the buyer and seller interact, describing how strategies and beliefs are defined as well as how information arrival works in our model. The timing of moves within each period in the game is summarized in Figure 1. Section 2.4 introduces our worst-case notion. To maintain focus, discussion of model assumptions is deferred to Section 5.3.

2.1. Environment

A seller (he) of a durable good (e.g., a house) interacts with a single buyer (she) in discrete time until a terminal date $T \leq \infty$.¹¹ The buyer can purchase the good at any time $t = 1, \dots, T$. She has unit demand and obtains utility v from purchasing, where v is drawn once at time 0 from a commonly known distribution F and remains fixed throughout the game. The distribution F has density f supported on a compact interval $[\underline{v}, \bar{v}] \subset \mathbb{R}_+$. The case $\underline{v} > 0$ is referred to as the “gap case,” while $\underline{v} = 0$ is the “no-gap case.” Both buyer and seller discount payoffs by a common factor $\delta \in [0, 1)$.

However, neither the buyer nor the seller observes the realization of v . Instead, the buyer may learn about v over time. Formally, a (static) *information structure* is a pair (S, π) , where S is a standard Borel space and $\pi : [\underline{v}, \bar{v}] \rightarrow \Delta(S)$ is a measurable mapping from valuations to

¹¹We handle the cases $T = \infty$ and $T < \infty$ separately.

probability distributions over signals. In the dynamic setting, the buyer receives signals according to such information structures in a history–dependent manner, as described below.

For expositional clarity, we use the terminology of “information structures” to refer to the static mapping from values to signals that the buyer observes in any period. We refer to the full, history-dependent sequence of information structures as an *information arrival process*.

2.2. Actions and Histories

Each period t begins with the seller choosing a price $p_t \in \mathbb{R}_+$. While the seller may randomize over prices, we assume that the realization of p_t is observed before the game continues.

After the price p_t is realized, the buyer observes a signal drawn according to some information structure, determined in that period. Throughout the paper, we assume that the information structure as well as the realized signals are observed only by the buyer, and *not* by the seller.

After observing the price for the given period as well as the new information described in the previous paragraph, the buyer then updates beliefs and decides whether to purchase or not. Formally, the buyer’s action in period t is denoted $a_t \in \{0, 1\}$, where $a_t = 1$ indicates purchase at price p_t and $a_t = 0$ indicates no purchase. If $a_t = 1$ or $t = T$, the game ends; otherwise, play proceeds to period $t + 1$.

The buyer and the seller face distinct information sets. Since the buyer’s only decision is whether to purchase, and since the game ends once she does, our definition of histories conditions on the event that the buyer has not yet purchased. We let H_S^t and H_B^t denote the sets of feasible seller and buyer histories, respectively, at date t . A seller history consists of past realized prices, so that $H_S^t = \{(p_1, \dots, p_{t-1})\}$. A buyer history at date t , by contrast, includes the period t price, as well as the new information the buyer receives that period, summarized by (π_t, s_t) . Therefore,

$$H_B^t = \{(p_1, \pi_1, s_1, \dots, p_t, \pi_t, s_t)\}.$$

In particular, H_B^t does *not* include information from future periods.

Although we have not yet specified how information arrival is determined, it is worth noting that the framework so far is standard. For example, the environment reduces to textbook bargaining with one-sided private information if π_1 revealed v to the buyer. The novelty lies in allowing the buyer to learn about v over time, and in specifying a solution concept which allows for endogenous interactions between this information and the seller’s strategy.

2.3. Defining Strategies and Beliefs

We now describe strategies and beliefs for the players. This step allows us to define our solution concept, which captures the seller's informationally robust (i.e., worst-case) objective. Throughout the paper, we take “informationally robust” to mean with respect to the worst-case information arrival process. As discussed in the introduction, part of our contribution lies in formulating the corresponding solution concept under limited commitment. To our knowledge, there is no consensus approach on how to do so.

Strategies. A *pricing strategy* is a function

$$\sigma : \bigcup_t H_S^t \rightarrow \Delta(\mathbb{R}_+),$$

so that for each seller information set, σ specifies a distribution over prices. A *price path* is a sequence $(p_1, p_2, \dots)$; we write $p^t = (p_1, \dots, p_t)$ for the history of prices up to period t .

A *buyer strategy* is a function

$$\alpha : \bigcup_t H_B^t \rightarrow \Delta(\{0, 1\}),$$

where, for each buyer information set, α specifies a probability distribution over $\{0, 1\}$: as mentioned, 0 denotes “not buying,” and 1 denotes “buying.”¹²

Information Arrival Process. Let H_N^t be the set of price histories up to and including time t , together with the set of all possible information structures and signal realizations before time t . An *information arrival process* is a function

$$\Pi : \bigcup_t H_N^t \rightarrow \Delta(\{(S, \pi)\}),$$

which assigns, to each element of H_N^t , a distribution over the (static) information structure for that period.

Belief System. Let

$$\mathcal{H} := \bigcup_{t, i \in \{S, B, N\}} H_i^t$$

¹²In the standard Coasian bargaining model with a continuum of types, mixed strategies for the buyer are unnecessary for equilibrium existence. However, with general information arrival, mixed strategies may be necessary.




Figure 1: Timing of moves and information sets within some period t .

Two histories h and h' are *non-contradictory* if they coincide whenever possible for them to do so. Given (σ, α) and Π as above, define

$$\mathbb{P}_{\sigma, \alpha, \Pi} : \mathcal{H} \rightarrow \Delta([\underline{v}, \bar{v}] \times \mathcal{H})$$

to be the probability distribution over (v, h') induced when starting from history $h \in \mathcal{H}$ and given a profile (σ, α, Π) . For every $h \in \mathcal{H}$, $\mathbb{P}_{\sigma, \alpha, \Pi}(h)$ is supported on histories h' that are non-contradictory with h .

A *belief system* for $j \in \{S, B, N\}$ is a function

$$\mu_j : \bigcup_t H_j^t \rightarrow \Delta([\underline{v}, \bar{v}] \times \mathcal{H}),$$

such that at each h_j^t , μ_j is supported on histories $h \in \mathcal{H}$ consistent with that information set. Let $\mu = (\mu_S, \mu_B, \mu_N)$.

A belief system *satisfies Bayes' rule where possible* if, for $t < s$ and h^t non-contradictory with h^s , $\mu(h^s)$ is obtained from $\mu(h^t)$ via Bayes' rule. Since the buyer observes the complete history of prices, information structures and signal realizations, their beliefs concern only the private type v . By contrast, the seller observes neither the information structure nor the signal realization; he observes only past prices. Accordingly, the seller's beliefs are defined over both v and the unobserved history of information structures and signals. For any given seller information set, these beliefs are degenerate on the observed prices.

2.4. Equilibrium Assumptions

We finally specify our solution concept. Fix an arbitrary triple (σ, α, Π) , the induced measure $\mathbb{P}_{\sigma, \alpha, \Pi}$, and a belief system μ . Let h_S^t, h_B^t, h_N^t denote representative elements of H_S^t, H_B^t, H_N^t . Let $F_i^t(\cdot)$ denote the prior $F(\cdot)$ conditional on h_i^t for $i \in \{S, B, N\}$.

The buyer's strategy is *sequentially rational* given μ if, for all h_B^t , the action prescribed by α maximizes the buyer's expected continuation payoff conditional on reaching h_B^t :

$$\mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha, \Pi}} \left[\sum_{\tau \geq t} \delta^{\tau-t} (v - p_\tau) \mathbf{1}_{\{\text{accept at } \tau\}} \mid h_B^t \right],$$

where τ is the induced stopping time.

If the buyer purchases at some time s at price p_s , then from the perspective of time $t < s$ the seller obtains payoff $\delta^{s-t} p_s$. The seller's pricing strategy is *sequentially rational* given μ if, at every h_S^t , the action prescribed by σ maximizes the seller's expected continuation payoff conditional on reaching h_S^t :

$$\mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha, \Pi}} \left[\sum_{k=t}^T \delta^{k-t} p_k \mathbf{1}_{\{a_k=1\}} \mid h_S^t \right], \quad (1)$$

where $a_k \in \{0, 1\}$ denotes the buyer's acceptance decision at time k .

Definition 1. Let (σ, α) be strategies for the seller and the buyer, let Π be an information arrival process, and let μ be a belief system. The quadruple $((\sigma, \alpha), \Pi, \mu)$ is an **equilibrium** if and only if:

- σ solves, at every h_S^t ,

$$\sup_{\hat{\sigma}} \mathbb{E}_{\mu, \mathbb{P}_{\hat{\sigma}, \alpha, \Pi}} \left[\sum_{k=t}^T \delta^{k-t} p_k \mathbf{1}_{\{a_k=1\}} \mid h_S^t \right]; \quad (2)$$

- α solves, at every h_B^t ,

$$\sup_{\hat{\alpha}} \mathbb{E}_{\mu, \mathbb{P}_{\sigma, \hat{\alpha}, \Pi}} \left[\sum_{\tau \geq t} \delta^{\tau-t} (v - p_\tau) \mathbf{1}_{\{\text{accept at } \tau\}} \mid h_B^t \right];$$

- Π solves, at every h_N^t ,

$$\inf_{\hat{\Pi}} \mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha, \hat{\Pi}}} \left[\sum_{k=t}^T \delta^{k-t} p_k \mathbf{1}_{\{a_k=1\}} \mid h_N^t \right]; \quad (3)$$

- μ is derived from strategies using Bayes' rule wherever possible; and
- The distribution of v induced by μ at time t is **consistent** with Bayesian updating via information structures and signal realizations,

$$v \mid (\pi_1, s_1), \dots, (\pi_t, s_t), \quad (4)$$

where (π_k, s_k) denotes the information structure chosen and signal realized in period k .

We call an information arrival process induced in such an equilibrium *sequentially worst case*. The key point is that the buyer's strategy depends only on h_B^t , rather than jointly on h_B^t and the entire process Π .

We briefly interpret the equilibrium concept. At every history, the seller holds a belief over the buyer's valuation induced by the belief system μ and chooses a strategy to maximize payoff against the worst-case information arrival process. To formalize this objective, the seller must conjecture both the information arrival process and the buyer's strategy. A consistency requirement is therefore imposed. The conjectured buyer's strategy must be rational, meaning that given the conjectured process and pricing strategy, the buyer's behavior is optimal.

In our model, the seller observes only whether the buyer purchases. Accordingly, this formulation can be understood as operating entirely through the seller's conjectures: at each history, the seller forms beliefs about both the information arrival process and the buyer's strategy. The conjectured pair must be mutually consistent: the information arrival process is worst case given the conjectured buyer strategy, while the buyer strategy is optimal given the conjectured pricing strategy and information arrival process. As usual in an equilibrium definition, deviations in one component are evaluated holding the remaining strategies fixed; in particular, a deviation in the information arrival process may change the buyer history h_B^t , at which the fixed strategy α prescribes the corresponding action.

Notice that the buyer's posterior over v must be generated by an actual sequence of information structures and realized signals; it cannot be an arbitrary off-path belief induced by prices or other strategic events. In particular, the buyer updates beliefs about v only through the information arrival process, both on and off the equilibrium path.¹³

2.5. Payoff Guarantee Interpretation

Although our paper differs from the robust mechanism design literature because the seller has limited commitment in our setting, the solution concept also admits a payoff guarantee interpretation. Consider the standard Coasian bargaining environment in which the buyer knows v at time 0. In that setting, the equilibrium outcome generates a price path $(p_1^*, p_2^*, \dots)$. The same price path arises as a *Nash* equilibrium if, instead of choosing a new price in each period, the seller selects a complete pricing strategy at time 0 with future prices hidden, so that the buyer observes the price in a given period only when that period arrives. Thus, the seller's sequential rationality

¹³This restriction is in the spirit of "no-signalling-what-you-don't-know" refinements (Fudenberg and Tirole, 1991). Otherwise, one could construct equilibria in which a deviation is deterred by the buyer adopting the belief that $v = 0$ with probability 1, even if this lies outside the support of F .

is equivalent to a formulation in which the seller commits to a pricing strategy but prices are revealed to the buyer period by period.

In this same spirit, consider the following normal form game, which we refer to as the *hidden-price-hidden-information normal form*:

- The seller's action set is the set of all randomized pricing strategies $\sigma : \cup_t H_S^t \rightarrow \Delta(\mathbb{R}_+)$.
- Nature's action set is the set of all history-dependent information arrival processes

$$\Pi : \bigcup_t H_N^t \rightarrow \Delta(\{(\pi, S)\}),$$

Nature's payoff is the negative of the seller's discounted profit, so Nature chooses Π to minimize the seller's expected payoff.

- After the seller and Nature simultaneously choose their actions, the buyer chooses α , given μ (with μ derived from Bayes' rule given σ and Π), for all h_B^t , to maximize the continuation payoff conditional on reaching h_B^t :

$$\mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha, \Pi}} \left[\sum_{\tau \geq t} \delta^{\tau-t} (v - p_\tau) \mathbf{1}_{\{\text{accept at } \tau\}} \mid h_B^t \right],$$

where τ is the induced stopping time.

Crucially, the buyer's strategy here depends only on h_B^t , the set of prices and information observed up to time t . We have the following:

Proposition 1. *The Nash equilibrium in the hidden-price-hidden information normal form induces the same equilibrium outcome as in the equilibrium of the main model.*

Proposition 1 provides a mapping between our worst-case formulation and those from static robustness settings, where optimal outcomes correspond to a saddle point: choices are optimal against the worst-case realizations from a given set of possibilities. While described as a Nash equilibrium outcome, as mentioned above, the implication is that the solution corresponds to a saddle-point of an objective for which no such literal interpretation is necessary.

If the buyer knows v ex ante, this normal-form game reduces to canonical Coasian bargaining. If the buyer's strategy is treated as exogenous, this formulation resembles a single-agent decision problem under a worst-case objective. The normal-form game can therefore be viewed as a canonical extension that combines these two cases.

3. SOLUTION OF THE MAIN MODEL

We now characterize both the equilibrium price path and the induced information arrival process under sequentially worst-case. Since equilibrium existence is not assumed a priori, the construction also establishes that our formulation yields a coherent definition of worst-case information arrival.

3.1. Arbitrary Finite Horizon

Equilibrium buyer behavior under sequentially worst-case information turns out to be characterized by a sequence of thresholds adapted to the seller's pricing strategy.

Definition 2. Fix any seller strategy σ . For any realized p_t , let $w_t(p_t, \sigma)$ be the unique value satisfying

$$w_t(p_t, \sigma) - p_t = \max_{\tau \geq t+1} \mathbb{E}[\delta^{\tau-t}(w_t(p_t, \sigma) - p_\tau) \mid p^t],$$

where τ ranges over all stopping times.

The **myopic threshold** information arrival process is defined as the process where, at any h_N^t , the buyer learns whether $v > y_t$, where y_t satisfies

$$w_t(p_t, \sigma) = \mathbb{E}_{v \sim F_N^t}[v \mid v \leq y_t]. \quad (5)$$

The terminology reflects the property that y_t minimizes the probability of purchase at time t , given the continuation strategy σ when $v \sim F_N^t$. This problem, in turn, is equivalent to a static Bayesian persuasion problem in which Nature (as Sender) seeks to induce the buyer (as Receiver) not to purchase. Thus, myopic refers to the *worst-case* objective used to determine the information structure each period, which may itself depend on the conjectured (future) price path. While Nature's objective is formally to minimize the seller's *discounted profit*, in our setting this coincides with minimizing the *period- t sale probability*.

Theorem 1. When $T < \infty$, the equilibrium price path is unique and deterministic (following possible randomization in period $t = 1$) and weakly decreasing over time. Every equilibrium induces the same buyer purchase behavior as the myopic-threshold information-arrival process with respect to the equilibrium price path.¹⁴

3.1.1. Illustrating Theorem 1 with $T = 2$ and Uniform Values

The general economic forces behind Theorem 1 can be seen in the simple case when there are only two periods to sell. We sketch the equilibrium construction for this case, taking $F = U[0, 2]$, and

¹⁴Nature may in principle provide more information than the threshold structure, provided this does not alter the buyer's behavior.

leave the formal details behind the claims made as part of this illustration to the general proof presented in Appendix A.

Given any discount factor δ , Theorem 1 yields the seller's equilibrium prices:

$$p_1^* = \frac{(2 - \delta)^2}{8 - 6\delta}, \quad p_2^* = \frac{2 - \delta}{8 - 6\delta}.$$

In each period $t \in \{1, 2\}$, the buyer learns whether $v > y_t^*$, where:

$$y_1^* = \frac{4 - 2\delta}{4 - 3\delta}, \quad y_2^* = \frac{4 - 2\delta}{8 - 6\delta}.$$

Buyers with $v > y_1^*$ purchase in period 1. Those with $y_1^* \geq v > y_2^*$ are indifferent between purchasing in period 1 or waiting until period 2, and in equilibrium purchase in period 2 (with ties broken against the seller). Buyers with $v \leq y_2^*$ are indifferent between purchasing in period 2 or never, and in equilibrium never purchase. Since $F(y) = \frac{y}{2}$ for the uniform distribution, the seller's equilibrium profit is:

$$\pi = p_1^* (1 - F(y_1^*)) + \delta p_2^* (F(y_1^*) - F(y_2^*)) = \boxed{\frac{(2 - \delta)^2}{4(4 - 3\delta)}}.$$

We now outline the key steps leading to this equilibrium.

Step One: Worst-Case Information is Threshold Information in the Second Period. For any second-period price p_2 , worst-case information solves an information design problem in the spirit of Kamenica and Gentzkow (2011): Nature chooses information to persuade the buyer not to buy. If $p_2 \geq \mathbb{E}_{F_N^1}[v]$, the buyer would not purchase even without any additional information, yielding zero profit to the seller. We therefore focus on the case $p_2 < \mathbb{E}_{F_N^1}[v]$ and p_2 lies above the lower bound of $\text{supp}(F_N^1)$ (as otherwise, no possible information structure would deter purchase).

Since the buyer has a binary action set and the state space is continuous, Kolotilin (2015) implies that the solution involves revealing whether $v > y_2$, where y_2 satisfies¹⁵

$$p_2 = \int_{\underline{v}}^{y_2} v f(v \mid v \leq y_2) dv. \tag{6}$$

Under this threshold structure, the seller's more-preferred action (purchase) occurs when $v > y_2$, while the less-preferred action (no purchase) occurs when $v \leq y_2$. The threshold is uniquely pinned down by the indifference condition: the buyer's expected valuation conditional on a no-

¹⁵The case in which the posterior distribution has atoms is measure zero and is ignored here.

purchase recommendation must equal p_2 . Attempting to reduce the purchase probability further would require raising the buyer's conditional expectation given a no-purchase signal, which would induce purchase and hence make the deviation infeasible.¹⁶

Step Two: Determining the Second-Period Price. Step One characterizes worst-case information for any p_2 . Suppose the first period information structure also takes the form of a threshold: the buyer purchases immediately if v exceeds this threshold, delaying otherwise.

Under this assumption, condition (6) and the fact that $F = U[0, 2]$ imply that the profit-minimizing second-period threshold is $y_2 = 2p_2$. Let $y_1(p_1)$ denote the first-period sequentially worst-case threshold. Because the buyer purchases whenever $v > 2p_2$, the seller chooses p_2 to maximize expected profit

$$p_2 \left(1 - \frac{2p_2}{y_1(p_1)} \right),$$

yielding the optimal second-period price

$$p_2(p_1) = \frac{y_1(p_1)}{4}.$$

Step Three: Determining the First-Period Indifference Point for the Buyer. In equilibrium, the function $p_2(p_1)$ specifies the second-period price for *any* on- or off-path p_1 . By Step One, a buyer who delays is indifferent in period 2 between purchasing and not purchasing. Consequently, her *expected* payoff at $t = 1$ would be unchanged even if she were to *always* purchase in the final period.

It follows that a buyer indifferent between immediate purchase and delay at $t = 1$ is also indifferent between purchasing in period 1 and purchasing in period 2. Given a signal s_1 , this indifference requires

$$\mathbb{E}[v \mid s_1] - p_1 = \delta(\mathbb{E}[v \mid s_1] - p_2(p_1)) \implies \mathbb{E}[v \mid s_1] = \underbrace{\frac{p_1 - \delta p_2(p_1)}{1 - \delta}}_{=: w_1(p_1)}.$$

This characterization holds *even off path*: in Step One, the solved worst-case information structure did not depend on whether the realized p_2 was chosen on path, and the same reasoning applies were the seller to deviate from the conjectured p_1 .

¹⁶Relative to standard information design, an additional technical issue is that our analysis requires assumptions about off-path buyer behavior. Nature can provide strict incentives to the buyer while increasing the seller's profit by only an arbitrarily small amount. The limiting profit must also be achieved in equilibrium; otherwise Nature would profitably deviate. Hence, the buyer breaks ties against the seller, even off path.

Step Four (Key Step): Finding the First-Period Information Structure. The problem of finding worst-case information in period 1 is now analogous to the period-2 problem in Step One. Seller profit is minimized if the buyer learns whether $v > 2w_1(p_1)$. The first-period threshold is:

$$y_1(p_1) = 2w_1(p_1).$$

The key observation is that the seller's payoff cannot be lowered beyond what this threshold strategy achieves. His optimal period-2 choice already ensures that a buyer delaying to period 2 is indifferent at the price $p_2(p_1)$. Therefore, the buyer is indifferent between purchasing in period 1 and delaying *if and only if* her expected valuation equals $w_1(p_1)$. Even following a deviation of Nature, this conclusion remains given the seller's last-period choice.

Thus, the first-period worst-case information problem is essentially identical to that in Step One, except that the relevant indifference point is now $w_1(p_1)$ instead of p_2 .

Step Five: Putting Everything Together to Find the Optimal First-Period Price. From Step Two, $p_2(p_1) = y_1(p_1)/4$, and from Step Four, $y_1(p_1) = 2w_1(p_1)$. Combining these yields

$$p_2(p_1) = \frac{w_1(p_1)}{2}.$$

Substituting into the definition of $w_1(p_1)$ from Step Three gives

$$w_1(p_1) = \frac{2p_1}{2 - \delta}.$$

Given p_1 , the seller's expected profit is

$$p_1 \left(1 - \frac{y_1(p_1)}{2} \right) + \delta p_2(p_1) \left(\frac{y_1(p_1)}{2} - \frac{y_2(p_1, p_2(p_1))}{2} \right).$$

Since $y_1(p_1) = 2w_1(p_1)$ and $y_2(p_1, p_2(p_1)) = 2p_2(p_1) = w_1(p_1)$, this simplifies to

$$p_1 \left(1 - \frac{2p_1}{2 - \delta} \right) + \delta \frac{p_1}{2 - \delta} \left(\frac{2p_1}{2 - \delta} - \frac{p_1}{2 - \delta} \right).$$

Maximizing over p_1 yields

$$p_1^* = \frac{(2 - \delta)^2}{8 - 6\delta},$$

as stated earlier. The equilibrium values y_1^* , y_2^* , and p_2^* then follow directly.

3.1.2. Discussion of the Solution

The central economic implication of this analysis is that the value at which the buyer is indifferent between immediate purchase and delay depends *only* on the conjectured price path. It does *not* depend on the possibility of future information. This point was illustrated in Step 4, making it the key step. Consequently, the prospect of additional information in the second period does not increase the probability of delay, thereby justifying the first-period's threshold information structure as worst-case.

By contrast, if it is common knowledge that the buyer were to learn v perfectly in period 2 for exogenous reasons, the first-period worst-case threshold would generally need to be higher. In that case, any buyer with $\mathbb{E}_{F_N^1}[v] \leq w_1(p_1)$ would *strictly* prefer to delay: knowing v exactly in period 2 allows her to avoid purchasing when $v < p_2$, delivering additional surplus. The seller's payoff could decrease if the first-period threshold were increased to induce more delay. However, this case does not arise under the equilibrium definition.

One of the main economic takeaways from the above analysis is that the sequentially worst-case formulation enables the seller to rely on simple pricing heuristics, avoiding much of the complexity involved in optimizing against *arbitrary* information arrival processes. Our analysis implies that sellers concerned about buyer-side information arrival can adopt strategies closely resembling those from the perfect-information (buyer side) benchmark. In the general case with myopic threshold information, the seller's discounted expected profit from time t onward (taking $w_0 = y^*(w_0) = \bar{v}$) is:

$$\sum_{s=t}^T \delta^{s-t} p_s \frac{F(y^*(w_{s-1})) - F(y^*(w_s))}{F(y^*(w_{t-1}))}.$$

By comparison, when the buyer knows v perfectly, the seller's discounted profit resembles this expression, with the only difference being that the purchase threshold in period $s \in \{t, \dots, T\}$ is w_s rather than $y^*(w_s)$. Thus, differences in the seller's objective relative to the known-value case arise entirely through the function $y^*(\cdot)$, which can be computed to primitives.

3.1.3. Additional Challenges in the Proof

The two-period case captures the basic intuition underlying Theorem 1. The same logic applies generally: for sequentially worst-case information, the buyer cannot expect to obtain payoff-relevant information if delaying, thus implying the solution is a myopic threshold information arrival process. However, extending this reasoning to the general case is substantially more involved. Indeed, our exercise *requires* solving for equilibrium strategies in a *three-player* game (seller, buyer, Nature) for which establishing general existence and uniqueness is non-trivial.

We mention two technical issues in our proof beyond those discussed for the $T = 2$ case:

(1) Handling arbitrary past information structures. In the informal construction, we solved for the second-period information structure *assuming* the first-period information structure were partitional. This property need not hold in equilibrium (as other information structures can induce buyer indifference following a recommendation to not buy). More significant, however, is that our general proof requires us to consider *arbitrary* (past) information structures, which cannot be ruled out a priori—that is, before we determine what the worst-case actually is. Showing that the intuition from Section 3.1.2 continues to hold in these general cases requires additional arguments.

(2) Allowing for on-path seller randomization. Our use of backward induction reasoning prevents us from imposing assumptions on the buyer’s *posterior* value distribution at any given time—in particular, we cannot make assumptions that would rule out “early” randomizations. Indeed, the posterior value distribution depends endogenously on the equilibrium information arrival process. Our proof accommodates the possibility that the seller might have randomized early, in anticipation of an exotic information structure arising later. While our theorem ultimately implies that on the equilibrium path seller randomization can occur only in the first period and is pinned down after that—consistent with the standard results in Fudenberg et al. (1985) and Gul et al. (1986)—ruling out randomization in later periods is a substantive part of the argument.

3.2. The Gap Case with $T = \infty$

We now turn to the infinite-horizon case. Although the main focus of the paper is on the formulation of the problem and the structure of the sequentially worst-case information arrival process, it is natural as a theoretical exercise to compare our results with the standard Coasian bargaining benchmark. The main complication is that backward induction no longer applies in a straightforward way, since the standard skimming property is absent from our model. Nevertheless, when $\underline{v} > 0$, the key Coasian pricing and selection forces emerge in equilibrium.

Readers familiar with the bargaining literature may associate the assumption $\underline{v} > 0$ with the conclusion that the market clears in finite time—i.e., that there exists a uniform bound $\hat{T}(\delta)$ such that, after any history, the buyer purchases with probability one by period $\hat{T}(\delta)$. This conclusion, however, requires additional regularity assumptions on the distribution of willingness-to-pay, such as Lipschitz continuity near the lower bound of its support (see Gul et al. (1986)). In our setting, such regularity conditions may fail for posterior distributions that arise under some information arrival processes, as we do not rule out any information structure a priori.

It is straightforward to verify that the myopic threshold information arrival process can still be sustained as *an* equilibrium strategy, by the one-shot deviation principle. The more subtle question is whether all sequentially worst-case information arrival processes must induce the

same buyer behavior as the myopic-threshold process. We show that this form of uniqueness *can* be recovered within a class of equilibria satisfying an intuitive monotonicity property.

Definition 3. An equilibrium $((\sigma, \alpha), \Pi, \mu)$ is a **monotone equilibrium** if and only if the buyer's strategy α is weak-Markov—depending only on the posterior belief F_B^t and the price p_t —and whenever F_1 strictly FOSD F_2 , the condition $0 \in \alpha_t(F_1, p_t)$ implies $\alpha_t(F_2, p_t) = 0$.

Monotonicity requires that if the buyer's perceived distribution of willingness-to-pay becomes *uniformly less* favorable, her willingness to delay cannot increase. It is immediate that when $T = \infty$, the equilibrium with the myopic information arrival process is monotone. Moreover, under the same regularity assumption as in Gul et al. (1986), we recover uniqueness.

Proposition 2. Suppose $T = \infty$, $\underline{v} > 0$, and F^{-1} is Lipschitz-continuous at 0. Then there exists an essentially unique monotone equilibrium. The equilibrium price path is unique and deterministic (following possible randomization in period $t = 1$). In any such equilibrium, every sequentially worst-case information arrival process induces the same buyer behavior as the myopic threshold information arrival process.

The proof shows monotonicity restores the property that the market clears in finite time. Once this property holds, sequentially worst-case information implies that future information cannot be used to induce delay. Consequently, the seller's value function at time t can be written as:¹⁷

$$V(y_{t-1}(p_{t-1})) = \max_{p_t} \left[p_t (F(y_{t-1}(p_{t-1})) - F(y_t(p_t))) + \delta V(y_t(p_t)) \right], \quad (7)$$

with $y_0 = \bar{v}$ and $y_t(\cdot)$ defined by (5).

Several structural properties of the equilibrium—such as the weak-Markov property¹⁸ and the absence of on-path seller randomization after the initial period—follow directly from the existence of this representation, by applying standard arguments from, for example, Fudenberg et al. (1985) and Ausubel and Deneckere (1989).

4. BEYOND SEQUENTIALLY WORST-CASE VIA SAFE SOLUTIONS

The sharp characterization of Theorem 1 relies on the sequential formulation, with the defining feature being that the buyer conditions only on the realized history. In the solution to the model,

¹⁷This representation is guaranteed only under sequentially worst-case information; it is therefore a *result*, not an assumption, that it characterizes price-setting behavior. The proofs of Proposition 2 and Theorem 1 do not assume such a value function *a priori*.

¹⁸We call an equilibrium profile *weak-Markov* if the buyer's acceptance decision depends only on F_B^t and p_t . This coincides with the standard definition when the buyer knows v .

the sequentially worst-case information arrival process consists of progressively lower partitional thresholds, inducing indifference in the absence of any further information. A natural robustness question concerns what happens if the buyer could instead condition on the *entire* information-arrival process at the outset, thus allowing future learning to directly influence her decision whether to delay today.

We ask whether this additional conditioning could lower the payoff guaranteed by the pricing strategy from Theorem 1. Formally, we fix the seller’s equilibrium pricing strategy but allow the buyer to condition her strategy on the complete information-arrival process. We call the pricing strategy safe if this enlargement cannot lower the seller’s payoff guarantee:

Definition 4. *An equilibrium pricing strategy is a **safe solution** if, at any on-path history, the seller’s equilibrium payoff equals the worst-case profit obtained from using the same pricing strategy against an arbitrary information arrival process, allowing the buyer’s strategy to depend not only on realized histories but also on the entire information arrival process—including the information structures specified at future histories.*

For sequentially worst-case information, the buyer’s strategy depends on the history h_B^t alone, rather than the information arrival process Π .¹⁹ Allowing the buyer’s strategy to depend on Π , as Definition 4 considers, might make delay more attractive if the buyer anticipates more favorable future information. Equivalently, allowing the complete process to shape current behavior gives Nature additional ability to use future information to influence the buyer’s stopping decision.

Of course, one possibility is that this future dependence coincides with the buyer’s conjecture under Theorem 1, in which case the buyer’s strategy would be unchanged. But in general, allowing the buyer and the seller to commonly observe the information arrival process could yield lower seller profit than the level identified in Theorem 1. Appendix C.1 uses the same two-period uniform environment as Section 3.1.1 to construct a commonly-known information process that induces an equilibrium with strictly lower seller profit than the benchmark in Theorem 1. The construction makes waiting more attractive by giving the buyer additional continuation surplus while preserving the seller’s incentives not to deviate.

The following property is sufficient for the price path in Theorem 1 to be safe.

Definition 5. *Let $y(w)$ satisfy $w = \mathbb{E}_{v \sim F}[v \mid v \leq y(w)]$. We say that a distribution F is threshold-ratio monotone if*

$$\frac{y(w)}{w} \text{ is weakly increasing in } w.$$

¹⁹At an on-path continuation history, past information is treated as sunk: the seller conditions on the history that has occurred rather than allowing Nature to re-optimize retrospectively over the information structures that generated it. Allowing such re-optimization changes the selection of buyers reaching the continuation history; see the longer working paper version for a discussion.

Recall that the myopic threshold information arrival process in Theorem 1 informs the buyer whether $v > y(w_t(p_t))$. Intuitively, threshold-ratio monotonicity ensures that increasing the threshold to induce more delay does not disproportionately raise the conditional expectation $\mathbb{E}[v \mid v \leq y]$. In other words, the expected value of buyers below the threshold in earlier periods changes less than the expected value in later periods. This guarantees that any reduction in early-period revenue is smaller than the corresponding gain from later-period sales, as our second main Theorem shows:²⁰

Theorem 2. *Suppose the value distribution is threshold-ratio monotone. Then the equilibrium pricing strategy in Theorem 1 and Proposition 2 is a safe solution: if the seller uses the strategy described there, no information arrival process can reduce the seller's expected payoff at any on-path history even if the buyer's strategy can condition on the process.*

Theorem 2 explicitly solves for the worst-case information arrival process when the buyer's strategy can condition on the process under the assumption of threshold-ratio monotonicity, and shows that the worst-case payoff is attained by the myopic-threshold process from Theorem 1. The proof relies on the fact that the equilibrium price path is deterministic (after the initial period), but otherwise places no restriction on the seller's strategy.

A useful preliminary observation is that, given a fixed price path, even if the buyer can condition her strategy on the whole information arrival process, the worst-case process is partitional in each period. Even so, substantial work remains because the worst-case thresholds may differ from those implied by the sequentially worst-case process. Recall that under sequentially worst-case information, whenever the buyer delays, she is indifferent between delaying and purchasing. Put differently, threshold information structures could still make the buyer *strictly* prefer to delay. Thus, solving for worst-case information involves a non-trivial choice of a threshold for each period, subject to the buyer's obedience constraints.

We address this by identifying a specific adjustment of the partition thresholds that lowers discounted profit whenever a threshold fails to induce exact indifference after a recommendation not to buy. While lowering the threshold increases sales in that period, we adjust the previous period's threshold to preserve obedience. In Appendix A, we verify that, under threshold-ratio monotonicity, this adjustment strictly reduces the seller's profit.

Theorem 2 should be interpreted as a robustness result for pricing strategies that are sustainable under limited commitment. The safe-solution criterion deliberately does *not* optimize payoff guarantees over arbitrary price paths, which need not remain optimal when future dates arrive. Rather, it asks whether the equilibrium pricing strategy continues to deliver its payoff guarantee

²⁰Appendix C.2 presents an example illustrating that the sequentially worst-case equilibrium outcome need not be a safe solution, so that some condition on the value distribution is needed for Theorem 2 to hold.

when the informational worst case is enlarged. The theorem shows that, under threshold-ratio monotonicity, it does.

While not every distribution satisfies threshold-ratio monotonicity, the condition is still quite permissive. The following sufficient condition illustrates this point.

Proposition 3. *For differentiable f , threshold-ratio monotonicity holds if $\frac{vf(v)}{F(v)}$ is decreasing in v .*

Proposition 3 implies that the class of threshold-ratio monotone distributions includes $F(v) = v^a$ for $a > 0$, exactly the case considered in Fuchs and Skrzypacz (2013), for instance. Threshold-ratio monotonicity also holds for uniform distributions supported on compact intervals of $\mathbb{R}^+$. Consequently, the price path described in Section 3.1.1 is safe.

Finally, if threshold-ratio monotonicity holds, the restriction to monotone equilibria in Proposition 2 can be substantially relaxed.

Definition 6. *An equilibrium $((\sigma, \alpha), \Pi, \mu)$ is a **deterministic equilibrium** if and only if the equilibrium pricing strategy is deterministic except possibly in the first period.*

The actual price path from the second period onward may depend on the realization of the first-period price. Equilibria in the known-values case under the Lipschitz condition satisfy this restriction (see Fudenberg et al. (1985) and Gul et al. (1986)).²¹

Corollary 1. *If, in addition to the conditions of Proposition 2, F is threshold-ratio monotone, then the equilibrium in Proposition 2 is the unique deterministic equilibrium.*

4.1. Payoff Guarantee Interpretation of Safe Solutions

We conclude by pointing out that safe solutions emerge naturally in the saddle-point formulation analogous to the one described in Section 2.5:

Definition 7. *A profile $(\sigma, \alpha(\cdot | \Pi), \Pi)$ is a **one-shot worst-case** if the information arrival process Π minimizes*

$$\mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha(\cdot | \Pi), \Pi}} \left[\sum_{k=1}^T \delta^{k-1} p_k \mathbf{1}_{\{a_k=1\}} \right], \quad (8)$$

i.e., the ex-ante expected seller's discounted payoff, and the buyer's strategy is optimal at every h_B^t given Π .

²¹Gul et al. (1986) also discuss why the equilibrium pricing strategy may not be deterministic in the initial period, hence our inclusion of this condition in the definition.

While this notion maintains the assumption from Section 2.5 that the price is revealed to the buyer period-by-period, the key difference is that it allows the information arrival process to be revealed to the buyer *at the very beginning*. Consider the normal-form game that arises under this modification, which we call the *hidden-price-observed-information normal form*:

- The seller's action set is the set of all randomized pricing strategies $\sigma : \cup_t H_S^t \rightarrow \Delta(\mathbb{R}_+)$.
- Nature's action set is the set of all history-dependent information arrival processes

$$\Pi : \bigcup_t H_N^t \rightarrow \Delta(\{(\pi, S)\}),$$

As before, Nature chooses Π to minimize the seller's expected payoff.

- The buyer chooses $\alpha(\cdot \mid \Pi)$ for all h_B^t , given μ (with μ derived from Bayes' rule given σ and Π), maximizing the continuation payoff conditional on reaching h_B^t :

$$\mathbb{E}_{\mu, \mathbb{P}_{\sigma, \alpha(\cdot \mid \Pi), \Pi}} \left[\sum_{\tau \geq t} \delta^{\tau-t} (v - p_\tau) \mathbf{1}_{\{\text{accept at } \tau\}} \mid h_B^t \right],$$

where τ is the induced stopping time.

The difference with the payoff-guarantee interpretation in Section 2.5 arises—as suggested by the name—since the choice of Π is observed by the buyer. However, we still assume simultaneity in the choice of pricing strategy and the choice of information arrival process. We have the following Proposition:

Proposition 4. *The Nash equilibrium in the hidden-price-observed-information normal form game (between the seller and Nature) induces a one-shot worst-case information arrival process. A sequentially worst-case equilibrium pricing strategy is a safe solution only if it can be sustained in some Nash equilibrium of the hidden-price-observed-information normal form game.*

Taken together, Propositions 1 and 4 shed light on both the notions of sequentially worst-case information arrival process and safe solutions—and their connection as articulated in Theorem 2. Proposition 1 shows that the sequentially worst-case benchmark arises from restricting the set of possibilities (buyer strategies in particular) over which the seller considers worst-case outcomes, while Proposition 4 shows that safe solutions arise from removing these restrictions.

5. ADDITIONAL RESULTS AND DISCUSSION

5.1. Indeterminacy for Exogenous Information Arrival

Part of the motivation for our approach is the observation that classical Bayesian approaches have tended to restrict attention to specific forms of information arrival. Here, we show that allowing general commonly-known and exogenously fixed information-arrival processes can generate substantial indeterminacy in equilibrium predictions. More precisely, such processes can sustain equilibrium outcomes that differ dramatically from the known-values benchmark.²²

Proposition 5. *Fix F and δ , and take $T = \infty$. Suppose the equilibrium outcome when the buyer knows v does not involve market clearing at $t = 1$. Then there exists an exogenous information arrival process admitting an equilibrium (between the seller and the buyer under this process) such that:*

- *The seller follows a constant price path.*
- *The seller's expected payoff is v^* , where v^* is any value strictly less than $\mathbb{E}_F[v]$ and strictly greater than the equilibrium payoff identified in Proposition 2.*
- *The market does not clear in any finite time (i.e., there is no $\hat{T}$ such that the buyer purchases before time $\hat{T}$ with probability one on-path).*

When $T < \infty$, there exists an exogenous information arrival process and an equilibrium satisfying the first two points as long as v^ is strictly greater than the equilibrium payoff identified in Theorem 1.*

This result stems from a general source of indeterminacy due to information arrival. With commonly known information arrival, *continuation information* can depend on current prices, since the process may specify different future information following different realized price histories. This provides another channel for continuation play to discipline seller deviations. Proposition 5 shows just how substantial this force can be: a continuum of seller payoffs can be sustained, constant-price equilibria can arise, and in the infinite-horizon case the market need not clear in finite time. These possibilities do not arise in the known-values gap case.²³

The proof of Proposition 5 illustrates this force with a surprisingly simple information arrival process: On path, the seller sets price equal to $\mathbb{E}_F[v]$, no information is provided if the seller adheres to the constant price path, and following a deviation by the seller, play switches to the sequentially worst-case continuation described in Theorem 1 and Proposition 2. Since the

²²This exercise is in the spirit of *robust predictions*; Liu (2022) performs such an exercise when the *seller* may obtain extra information about the buyer's value, showing that a rich set of payoffs may emerge in the frequent-offer limit.

²³The combination with a constant price path is also distinct, although constant prices would arise in the degenerate case where the market clears at time 1.

buyer's expected payoff is 0 at every time, this construction supports an equilibrium in which the buyer randomizes with a probability that induces the seller to maintain the constant price path rather than deviating. By disciplining deviations, the prospect of a change in the informational environment can therefore sustain constant price paths and indeterminate buyer behavior, in contrast to the no-learning case. Our approach, by contrast, restores simple selection and pricing forces underlying Coasian dynamics even when information arrival is present. Proposition 5 shows that this recovery is nontrivial: at its extreme, unrestricted information arrival can lead to a complete breakdown of these predictions.²⁴

5.2. Discussion of the No-Gap Case

While our main analysis focused on either the finite-horizon or gap cases, many of our main insights also apply to the no-gap case:

Proposition 6. *Suppose $T = \infty$ and $\underline{v} = 0$. An equilibrium exists in which the worst case information arrival process is a myopic threshold information arrival process.*

Given that it is typically impossible to impose finite-time market clearing when $\underline{v} = 0$, monotonicity alone does not suffice to rule out non-myopic-threshold information arrival processes. However, the threshold-ratio monotonicity condition on F does imply that this outcome belongs to the same class of equilibria, with possible multiplicity:

Proposition 7. *Suppose $T = \infty$ and $\underline{v} = 0$. If F is threshold-ratio monotone, then in any deterministic equilibrium, the worst case information arrival process is equivalent to myopic threshold information arrival process.*

This proposition highlights the relevance of myopic-threshold information arrival process even in the no-gap case. Note that in the no-gap case we no longer obtain unique equilibrium price paths, just as in the known-values case—identical constructions such as in Ausubel and Deneckere (1989) apply once we restrict to myopic threshold information arrival process.

5.3. Discussion of Model Assumptions

Dynamic Consistency Our worst-case objective reflects a seller unwilling to specify a prior over buyer-side information arrival. While a seller might form a prior over the buyer's true

²⁴This conclusion is not a feature of every formulation of the worst-case objective. Our longer working paper version shows that other formulations may fail to inherit this intuition. For example, naiveté may lead the seller not to sell at all if she expects the buyer to have strong delay incentives.

willingness-to-pay from market data using transaction prices, he may be less familiar with up-to-date information sources. In many settings, those sources are intangible, difficult to observe, and evolving, making it demanding to form a well-defined prior over information arrival.

Dynamic consistency holds in our framework because the information arrival process the seller anticipates is precisely the one that materializes both on-path and after his own deviations. The longer working paper version of this paper presents a detailed exploration of alternative formulations. These alternatives also illustrate why the formulation matters for tractability and interpretation. Other formulations need not preserve the simple backward-inductive structure of sequentially worst-case information and can generate more complex characterizations.

Accommodating dynamic inconsistency in our exercise is more involved than in the single-agent case. Epstein and Schneider (2003) propose a “rectangularity” condition on the set of priors that characterizes when the maxmin decision rule is dynamically consistent, while also providing a procedure for constructing a set of priors to represent a decision maker’s ex-ante preference.²⁵ Fabbri and Ortoleva (2026) axiomatizes another model of belief updating under maxmin preferences, interpreting the problem as a game against Nature in which equilibrium behavior is characterized by saddle points. The paper highlights the impossibility of accommodating Ellsberg-style violations of expected utility under both consequentialism and dynamic consistency, and describes a sense in which rectangularity can be understood as arising from a particular assumption on the timing of Nature’s moves.

Our analysis differs because it focuses on a strategic setting. Formulating the seller’s problem requires specifying how the seller anticipates the buyer’s strategy over time, and how the buyer forms expectations about the seller’s behavior. Dynamic consistency in our model stems from how we formulate the *buyer’s* problem: the buyer’s strategy at a given history cannot depend on future realized information choices. Our formulation induces a recursive structure that captures the seller’s backward-inductive reasoning. One interpretation is simply that the seller anticipates that behavior arises in this way. Another is that the extensive form is such that information obtained at one date does not influence what information might become available later. Under this interpretation, allowing current behavior to depend on unrealized future information enlarges the conjectured environment beyond the dependence permitted by the extensive form.

Other Assumptions We briefly discuss other aspects of our model which are more familiar from other related work. First, we have in mind situations where the buyer can consult any source that comes to mind costlessly. The restrictions on what the buyer knows about v reflect limits on what the buyer has access to at any time, rather than an effort choice. We share this assumption

²⁵ Another approach to accommodating ambiguity aversion is introduced in Klibanoff et al. (2009), which develops a dynamically consistent version of smooth ambiguity preferences allowing for a richer separation between ambiguity and ambiguity attitudes than maxmin. Hanany et al. (2020) extends this framework to games.

about how information is generated with much of the informational robustness literature (e.g., Du (2018); Brooks and Du (2021, 2023); Deb and Roesler (2023), among others).

A general issue for robust objectives concerns timing. We have in mind situations where buyers have some time to respond to offers, or where new offers are made very shortly after rejection (so that the buyer can learn while considering the offer). While the seller may randomize, the seller’s action within a period (i.e., the posted price) could influence the information available to the buyer. Indeed, many papers have noted channels through which information can depend on price in practice—see, for instance, Xu and Yang (2022); Liu et al. (2023); Ichihashi and Smolin (2023). Ke and Zhang (2020) provide decision-theoretic foundations for the assumption that the worst case depends on realized choices, an assumption also adopted in other work (e.g., Carroll (2015); Bergemann et al. (2017); Chen (2023); Guo and Shmaya (2023); Malladi (2023)).

For us, allowing the worst case to condition on realized seller choices is more compelling for two reasons. First, allowing information to depend on the randomization but not on its realization evokes commitment, since it requires a seller who observes the outcome of the randomization not to reconsider (for instance, after observing the price draw but before posting it). Our interest, however, is in cases where no commitment ability is present. The inability to commit to a randomization has traditionally been used to justify restricting attention to deterministic mechanisms more generally (see Laffont and Martimort, 2002, p. 67). Second, assuming that information depends on the seller’s strategy and not realized choices implies that information is unchanged if the seller *deviates* from a prescribed strategy. Yet deviations may well be deterred by the prospect of influencing that period’s informational environment, so allowing some price dependence seems natural. As we see no universally plausible a priori restriction, we therefore leave this possibility unconstrained.

5.4. Concluding Remarks

We considered a seller who repeatedly reoptimizes prices against a worst-case information arrival process, obtaining a dynamically consistent objective and a sharp characterization of equilibrium outcomes. This formulation addresses conceptual issues that arise when commitment is relaxed under a worst-case objective. Durable goods pricing is a natural first setting to study these issues, as the buyer’s problem reduces solely to the choice of *when* to purchase. Sequentially worst-case information restores the selection logic underlying dynamic pricing without imposing parametric structure on buyer learning; this finding should not be taken for granted. As we showed, allowing general information arrival can destroy sharp equilibrium predictions. For a broad class of value distributions, the equilibrium pricing strategy retains its payoff guarantee even when the worst case need not be sequential. We hope the benchmark we articulate inspires future work to deliver concrete predictions when buyer learning is significant yet difficult to model.

References

- Avidit Acharya and Juan Ortner. 2017. Progressive Learning. *Econometrica* 85, 6 (2017), 1965–1990.
- Robert Ainsworth, Rajeev Dehejia, Cristian Pop-Eleches, and Miguel Urquiola. 2023. Why Do Households Leave School Value Added on the Table? The Roles of Information and Preferences. *American Economic Review* 113, 4 (April 2023), 1049–1082.
- Nabil Al-Najjar and Jonathan Weinstein. 2009. The ambiguity aversion literature: A critical assessment. *Economics and Philosophy* 25, 3 (November 2009), 249–284.
- Sarah Auster, Yeon-Koo Che, and Konrad Mierendorff. 2022. Prolonged Learning and Hasty Stopping: The Wald Problem Under Ambiguity. *Working Paper* (2022).
- Lawrence M. Ausubel and Raymond Deneckere. 1989. Reputation in Bargaining and Durable Goods Monopoly. *Econometrica* 57, 3 (1989), 511–531.
- Ian Ball and Deniz Kattwinkel. 2025. Robust Robustness. *Unpublished manuscript* (2025). Working paper.
- Pierpaolo Battigalli, Emiliano Catonini, Giacomo Lanzani, and Massimo Marinacci. 2019a. Ambiguity attitudes and self-confirming equilibrium in sequential games. *Games and Economic Behavior* 115 (2019), 1–29.
- Pierpaolo Battigalli, Alejandro Francetich, Giacomo Lanzani, and Massimo Marinacci. 2019b. Learning and self-confirming long-run biases. *Journal of Economic Theory* 183 (2019), 740–785.
- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2017. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica* 85, 1 (2017), 107–143.
- Peter Bergman, Eric Chan, and Adam Kapor. 2020. Housing Search Frictions: Evidence from Detailed Search Data and a Field Experiment. *Working Paper* (2020).
- Subir Bose and Arun Daripa. 2009. Dynamic Mechanism and Surplus Extraction Under Ambiguity. *Journal of Economic Theory* 144 (2009), 2084–2115.
- Subir Bose, Emre Ozdenoren, and Andreas Pape. 2006. Optimal Auctions with Ambiguity. *Theoretical Economics* 1 (2006), 411–438.
- Subir Bose and Ludovic Renou. 2014. Mechanism Design with Ambiguous Communication Devices. *Econometrica* 82 (2014), 1853–1872.

- Benjamin Brooks and Songzi Du. 2021. Optimal Auction Design with Common Values: An Informationally Robust Approach. *Econometrica* 89, 3 (2021), 1313–1360.
- Benjamin Brooks and Songzi Du. 2023. On the Structure of Informationally Robust Optimal Mechanisms. *Working Paper, University of Chicago and University of California, San Diego* (2023).
- Gabriel Carroll. 2015. Robustness and linear contracts. *American Economic Review* 105, 2 (2015), 536–563.
- Gabriel Carroll. 2017. Robustness and Separation in Multidimensional Screening. *Econometrica* 85, 2 (2017), 453–488.
- Sylvain Chassang. 2013. Calibrated Incentive Contracts. *Econometrica* 81, 5 (September 2013), 1935–1971.
- Yeon-Koo Che and Weijie Zhong. 2022. Robustly-Optimal Mechanism for Selling Multiple Goods. *Working Paper* (2022).
- Jaden Chen. 2023. Sequential Learning under Informational Ambiguity. *Working Paper* (2023).
- Ronald Coase. 1972. Durability and monopoly. *The Journal of Law and Economics* 15, 1 (1972), 143–149.
- Rahul Deb and Anne-Katrin Roesler. 2023. Multi-dimensional screening: Buyer-Optimal Learning and Informational Robustness. *Review of Economic Studies* Forthcoming (2023).
- Alfredo Di Tillio, Nenad Kos, and Matthias Messner. 2017. The Design of Ambiguous Mechanisms. *Review of Economic Studies* 84 (2017), 237–276.
- Songzi Du. 2018. Robust Mechanisms Under Common Valuation. *Econometrica* 86, 5 (2018), 1569–1588.
- Jetlir Duraj. 2020. Bargaining with Endogenous Learning. *Working Paper, University of Pittsburgh* (2020).
- Théo Durandard, Udayan Vaidya, and Boli Xu. 2024. Robust Contracting for Sequential Search. (2024). Working paper.
- Larry Epstein and Martin Schneider. 2003. Recursive Multiple Priors. *Journal of Economic Theory* 113 (2003), 1–31.

- Francesco Fabbri and Pietro Ortoleva. 2026. Updating against Nature. (2026).
- Daryl Fairweather, Matthew E. Kahn, Robert Metcalfe, and Sebastian Sandoval-Olascoaga. 2023. The Impact of Climate Risk Disclosure on Housing Search and Buying Dynamics: Evidence from a Nationwide Field Experiment with Redfin. *Working Paper* (2023).
- William Fuchs and Andrzej Skrzypacz. 2013. Bargaining with Deadlines and Private Information. *American Economic Journal: Microeconomics* 5, 4 (2013), 219–243.
- Drew Fudenberg, David Levine, and Jean Tirole. 1985. Infinite-horizon models of bargaining with one-sided incomplete information. In *Game-Theoretic Models of Bargaining*, Alvin Roth (Ed.). Cambridge University Press, Chapter 5, 73–98.
- Drew Fudenberg and Jean Tirole. 1991. Perfect Bayesian Equilibrium and Sequential Equilibrium. *Journal of Economic Theory* 53, 2 (April 1991), 236–260.
- Faruk Gul, Hugo Sonnenschein, and Robert Wilson. 1986. Foundations of Dynamic Monopoly and the Coase Conjecture. *Journal of Economic Theory* 39 (1986), 155–190.
- Yingni Guo and Eran Shmaya. 2023. Robust Monopoly Regulation. *Working Paper* (2023).
- Lu Han and William C. Strange. 2015. The Microstructure of Housing Markets: Search, Bargaining, and Brokerage. In *Handbook of Regional And Urban Economics*, Gilles Duranton, J. Vernon Henderson, and William C. Strange (Eds.). Vol. 5. Elsevier, Chapter 13, 813–886.
- Eran Hanany, Peter Klibanoff, and Sujoy Mukerji. 2020. Incomplete Information Games with Ambiguity Averse Players. *American Economic Journal: Microeconomics* 12, 2 (May 2020), 135–187.
- Wei He and Jiangtao Li. 2022. Correlation-Robust Action Design. *Journal of Economic Theory* 200 (2022), 105403.
- Shota Ichihashi and Alex Smolin. 2023. Buyer-Optimal Algorithmic Consumption. *Working Paper* (2023).
- Ashwin Kambhampati. 2025. Proper Robustness and the Efficiency of Monopoly Screening. *Unpublished manuscript* (2025).
- Emir Kamenica and Matthew Gentzkow. 2011. Bayesian Persuasion. *American Economic Review* 101, 6 (October 2011), 2590–2615.

- Shaowei Ke and Qi Zhang. 2020. Randomization and Ambiguity Aversion. *Econometrica* 88, 3 (2020), 1159–1195.
- Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. 2009. Recursive Smooth Ambiguity Preferences. *Journal of Economic Theory* 144, 3 (May 2009), 930–976.
- Andrew Koh and Sivakorn Sanguanmoo. 2025. Robust Technology Regulation. *Unpublished manuscript* (2025).
- Anton Kolotilin. 2015. Experimental design to persuade. *Games and Economic Behavior* 90 (2015), 215–226.
- Jean-Jacques Laffont and David Martimort. 2002. *Theory of Incentives*. Princeton University Press.
- Tuomas Laiho and Julia Salmi. 2020. Coasian Dynamics and Endogenous Learning. *Working Paper, University of Oslo and University of Copenhagen* (2020).
- Jonathan Libgober and Xiaosheng Mu. 2021. Informational Robustness in Intertemporal Pricing. *Review of Economic Studies* 88, 3 (2021), 1224–1252.
- Heng Liu. 2022. Robust Predictions in Coasian Bargaining. *American Economic Review: Insights* 4, 2 (2022), 209–222.
- Tingjun Liu, Dan Bernhardt, and Odilon Camara. 2023. Extracting surplus by walking away from acquiring information. *Working Paper* (2023).
- Niccolo Lomys. 2018. Learning while Bargaining: Experimentation and Coasean Dynamics. *Working Paper, Toulouse School of Economics* (2018).
- Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. 2020. Uncertainty in Mechanism Design. *Working Paper, Duke University, University of Pittsburgh and University of California, Berkeley* (2020).
- Suraj Malladi. 2023. Searching in the Dark and Learning Where to Look. *Working Paper* (2023).
- Paul Milgrom and Chris Shannon. 1994. Monotone Comparative Statics. *Econometrica* 62, 1 (1994), 157–180.
- Juan Ortner. 2017. Durable goods monopoly with stochastic costs. *Theoretical Economics* 12, 2 (2017), 817–861.

- Juan Ortner. 2023. Bargaining with Evolving Private Information. *Theoretical Economics* Forthcoming (2023).
- Alessandro Pavan. 2017. Dynamic Mechanism Design: Robustness and Endogenous Types. In *Advances in Economics and Econometrics*. Vol. 1. Cambridge University Press, 1–62.
- Antonio Penta. 2015. Robust Dynamic Implementation. *Journal of Economic Theory* 160 (December 2015), 280–316.
- Marciano Siniscalchi. 2009. Two out of three ain't bad: a comment on 'The ambiguity aversion literature: A critical assessment'. *Economics and Philosophy* 25, 3 (November 2009), 335–356.
- Jean Tirole. 2016. From Bottom of the Barrel to Cream of the Crop: Sequential Screening With Positive Selection. *Econometrica* 84, 4 (2016), 1291–1343.
- Alexander Wolitzky. 2016. Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade. *Theoretical Economics* 11 (2016), 971–1004.
- Wenji Xu and Kai Hao Yang. 2022. Informational Intermediation, Market Feedback, and Welfare Losses. *Working Paper* (2022).

A. PROOFS OF MAIN THEOREMS

Throughout this appendix, we use the common expositional device that information structures are chosen by some player, referred to as “Nature.” This simplifies the presentation of our solution concept.

Proof of Theorem 1.

Step 1. Reduction to Binary Recommendations. We first show that, without loss of generality, Nature can be assumed to provide only *binary* recommendations to the buyer in each period—“buy” or “wait.” This is the familiar “*recommendation principle*” from the information design literature.

Lemma 1. *Given equilibrium strategies (σ, α, Π) and the corresponding belief system μ , there exist equilibrium strategies $(\tilde{\sigma}, \tilde{\alpha}, \tilde{\Pi})$ and a belief system $\tilde{\mu}$ such that $\tilde{\Pi}$ provides only two signals (buy or wait) at every history h_N^t , and the induced equilibrium outcome is identical to that under $(\sigma, \alpha, \Pi, \mu)$.*

Proof. Let $(\pi_t, S_t) = \Pi(h_N^t)$ denote the information structure chosen by Nature at history h_N^t . With a slight abuse of notation, write $\alpha(h_B^t, s_t)$ for the probability that the buyer purchases after observing $s_t \in S_t$ in period t .

We first show that it is without loss to assume the buyer does not randomize. If $\alpha(h_B^t, s_t) \in (0, 1)$ for some s_t , Nature can “refine” the signal to implement the buyer’s mixed action deterministically. Specifically, construct a new information structure $(\tilde{\pi}_t, \tilde{S}_t)$ that, conditional on s_t , sends:

$$\bar{s}_t \quad \text{with probability} \quad \alpha(h_B^t, s_t), \quad \text{and} \quad \underline{s}_t \quad \text{otherwise,}$$

where the buyer always buys after $\bar{s}_t$ and never buys after $\underline{s}_t$. Let $\tilde{\alpha}$ coincide with α except that it is now deterministic. This transformation preserves all outcomes and incentives. Applying this to each (t, s_t) , we may assume $\alpha(\cdot)$ is deterministic.

With α deterministic, all signals on which the buyer buys can be pooled into a single “buy” signal $\bar{s}_t$, and all others into a single “wait” signal $\underline{s}_t$. Formally, replace π_t by

$$\tilde{\pi}_t : [\underline{v}, \bar{v}] \rightarrow \Delta(\{\underline{s}_t, \bar{s}_t\}),$$

where $\bar{s}_t$ is sent if and only if the original α prescribes purchase. To maintain any further distinctions lost by pooling at period t which could influence behavior in later periods, Nature can push the original signals forward at period t to period $t + 1$ if needed. More formally:

$$\tilde{\Pi}(h_N^t) = (\tilde{\pi}_t, \{\underline{s}_t, \bar{s}_t\}),$$

and after $(h_N^t, \tilde{\pi}_t, \tilde{s}_t, p_{t+1})$, Nature delivers the continuation information $(\pi_t, S_t) \cup \Pi(h_N^t, \pi_t, s_t, p_{t+1})$.

This construction changes neither the buyer's nor the seller's behavior, and the induced outcome is identical to that under $(\sigma, \alpha, \Pi, \mu)$. By induction on t , we conclude that Nature can be restricted, without loss, to binary recommendations at every history. $\square$

Step 2. Notation and Preliminaries. Given Lemma 1, we may assume without loss that Nature provides binary recommendations in every period. We now introduce notation used throughout the remainder of the proof.

Histories and posterior bounds. Fix a seller history h_S^t , and let $\underline{v}_S^t$ be the infimum of the support of F_S^t , the posterior distribution of v conditional on h_S^t .

Unconditional CDFs. For convenience, we express all Bellman iterations in terms of *unconditional* probabilities. Let

$$G_S^t(\cdot)$$

denote the CDF associated with a general Radon measure on $[\underline{v}, \bar{v}]$. Note that $G_S^t(\bar{v})$ need not equal 1; rather, it can be scaled to $F_S^t(\cdot)$ divided by the remaining buyer mass.²⁶

Threshold information structures. Given a history h_S^t and a scalar x , define

$$y_S^t(x) := \inf \{y \geq \underline{v}_S^t : x \leq \mathbb{E}_{F_S^t}[v \mid v \leq y]\}. \quad (9)$$

Intuitively, $y_S^t(x)$ is the threshold type below which the buyer's expected valuation is at most x .

Since F has a well-defined density, and—by Lemma 1—we may restrict attention to binary recommendations, F_S^t is continuous everywhere except possibly at histories where the buyer is certain to purchase. In particular, in (9) the inequality either holds with equality or yields $y_S^t(x) = \underline{v}_S^t$. Moreover, $y_S^t(\cdot)$ is continuous in x .

Buyer cutoff. Let $w_t = w_t(h_B^t)$ be the buyer expected valuation indifferent between buying immediately at t and delaying. This is purely for expositional convenience: at this point we have not established the skimming property or characterized the equilibrium. In principle, a low-posterior buyer could expect future information that increases her valuation.

Seller best-response correspondences. Define:

$$M_t(G_S^t) = \arg \max_{p_t} V^t(G_S^t(\cdot))$$

²⁶Formally, we work in the space $\text{Radon}([\underline{v}, \bar{v}])$ with the weak-* topology.

and

$$m_S^t(b) = \arg \max_{p_{t+1}} V^{t+1}(G_S^{b,t}(\cdot)),$$

where $G_S^{b,t}(\cdot)$ is the CDF obtained by truncating $G_S^t(\cdot)$ at b , and $V^t(\cdot)$ is the seller's expected payoff in period t .²⁷

Step 3. Backward induction We now proceed by backward induction. The general strategy is as follows: since, along the equilibrium path, the buyer's ex-ante value of future information is zero, Nature's best response in any given period is to *minimize* the probability of purchase in that period, as the equilibrium price path is decreasing.

Nature's last-period problem

Consider $t = T$.

Case 1: Market already cleared. If the market has already cleared, any strategy is optimal for Nature; in particular, she can provide a trivial "no information" structure.

Case 2: Market not cleared. Let F_S^T denote the posterior distribution of v among remaining buyers, with unconditional CDF G_S^T .

If $p_T \leq \underline{v}_S^T$, the buyer must purchase with probability 1 in equilibrium, regardless of Nature's choice.²⁸ In this case, the claim that Nature uses a threshold information structure giving the same buyer payoff as no information is trivially satisfied.

If instead $p_T > \underline{v}_S^T$, Step One of Section 3.1.1 implies the worst-case information structure is a *threshold* with cutoff $y_S^T(p_T)$. Under such a structure, one of two things happens:

- (a) The buyer is indifferent between purchasing and not when recommended to wait, or
- (b) Nature's choice has no effect on the buyer's action (e.g., if p_T is too high).

In either case, the buyer's expected payoff matches that under no new information. By sequential rationality, Nature's equilibrium last-period strategy at any history h_N^T is therefore a threshold rule. For any $p_T > \underline{v}_S^T$, Nature can approximate the worst-case profit while giving the buyer a *strict* incentive to buy by, e.g., revealing whether $v \leq y_S^T(p_T) - \varepsilon$ for arbitrarily small $\varepsilon > 0$. Therefore, in equilibrium, the buyer must break indifference *against* the seller; otherwise Nature could profitably deviate.

²⁷Since uniqueness is not yet established, $V^t(\cdot)$ need not be uniquely defined, and the weak-Markov property need not hold; the continuation value may depend on both the distribution and the full history. We therefore reference a particular equilibrium throughout.

²⁸This is a standard argument in the Coasian bargaining literature; see Fudenberg et al. (1985).

Seller's last-period problem

Consider the history p^{T-1} .

Case 1: Market should have cleared but has not. If under (Π, α) (where Nature and the buyer follow the equilibrium strategies, but possibly allowing for seller deviations), the market should have cleared but does not, then this history is off-path. We may assign arbitrary beliefs and prices to deter such deviations.²⁹

Case 2: Market not cleared on path. If the market should not have cleared and indeed has not, then the only equilibrium-consistent belief for the seller is that Nature and the buyer have followed (Π, α) . Thus the remaining posterior is F_S^T with unconditional CDF G_S^T .

For any p_T in the support of σ_T we must have:

$$p_T \in M_T(G_S^T) = \arg \max_{p_T} V^T(G_S^T(\cdot)) = \arg \max_{p_T} p_T (G_S^T(\bar{v}) - G_S^T(y_S^T(p_T))) .$$

Since both $G_S^T(\cdot)$ and $y_S^T(\cdot)$ are continuous, the objective is jointly continuous in (G_S^T, p_T) . By Berge's Maximum Theorem,

$$V^T(G_S^T) = \max_{p_T} p_T (G_S^T(\bar{v}) - G_S^T(y_S^T(p_T)))$$

is continuous in G_S^T ,³⁰ and $M_T(G_S^T)$ is non-empty and compact. This establishes the seller's optimal response.

Furthermore, the period- $T - 1$ indifference condition is:

$$w_{T-1} - p_{T-1} = \mathbb{E}_{\sigma_T} [\delta (w_{T-1} - p_T)] ,$$

since at $t = T - 1$ there is no ex-ante informational value from delaying. This pins down w_{T-1} uniquely, given anticipated future play, and implies:

$$p_{T-1} \geq \mathbb{E}_{\sigma_T} [p_T] .$$

The inductive step for Nature. Consider $t = k < T$, assuming the market has not yet cleared, and let F_S^k be the posterior distribution of v among remaining buyers.

Case 1: $p_k \leq \underline{v}_S^k$. If the seller posts $p_k \leq \underline{v}_S^k$, the buyer purchases with probability 1 in equilibrium, regardless of Nature's choice.

²⁹For example, the seller could post $\bar{v}$.

³⁰If $F_n \rightarrow F$ in the weak-* topology, then $F_n(x) \rightarrow F(x)$ only at all continuity points of $F(\cdot)$. Since we restrict to continuous CDFs, this technical issue does not arise.

Case 2: $p_k > \underline{v}_S^k$. By the induction hypothesis, Nature's continuation strategy from $t = k + 1$ onward employs thresholds $\{w_{k+s}\}_{s \geq 1}$ that yield the same expected payoff to the buyer as if no further information were provided. The unique cutoff w_k is therefore determined by:

$$w_k - p_k = \mathbb{E}_{\sigma_{k+1}} [\delta(w_k - p_{k+1})].$$

Let Nature's binary information structure be $(\pi_k, \{\bar{s}_k, \underline{s}_k\})$. If $\mathbb{E}_{F_S^k}[v \mid \pi_k, s_k] > w_k$, the buyer strictly prefers to purchase at $t = k$.

We claim: if $\bar{s}_k$ occurs with positive probability, then $\mathbb{E}[v \mid \underline{s}_k] = w_k$. Suppose instead $\mathbb{E}[v \mid \underline{s}_k] < w_k$. Modify π_k to $(\tilde{\pi}_k, \{\bar{s}_k, \underline{s}_k\})$ by reassigning $\bar{s}_k$ to $\underline{s}_k$ with probability $\varepsilon > 0$. For sufficiently small ε , $\mathbb{E}[v \mid \tilde{\pi}_k, \underline{s}_k] < w_k$, so the buyer waits after $\underline{s}_k$. At $t = k + 1$, Nature can split $\pi_k^{-1}(\underline{s}_k)$ from the new types $\tilde{\pi}_k^{-1}(\underline{s}_k) \setminus \pi_k^{-1}(\underline{s}_k)$, and revert to the original continuation strategy for the first group. The second group no longer buys at $t = k$; since the price path is decreasing by induction, the profit from them is at most $\delta \mathbb{E}_{\sigma_{k+1}}[p_{k+1}]$. Because

$$p_k \geq \mathbb{E}_{\sigma_{k+1}}[p_{k+1}] > \delta \mathbb{E}_{\sigma_{k+1}}[p_{k+1}],$$

this deviation (not detected by the seller) is strictly profitable for Nature — a contradiction. We also claim $\underline{s}_k$ corresponds exactly to buyers with $v \leq y_S^k(w_k)$. If not, there exist $v' > v''$ with $v' \in \pi_k^{-1}(\underline{s}_k)$ and $v'' \in \pi_k^{-1}(\bar{s}_k)$. Swapping v' and v'' leaves $\mathbb{E}[v \mid \bar{s}_k]$ above w_k (so buyer actions unchanged) but lowers $\mathbb{E}[v \mid \underline{s}_k]$ strictly below w_k . This reduces to the previous case and Nature can further reduce profit when the seller's strategy is fixed — again contradicting optimality.

Fixed-point condition. Because Nature can lower the threshold slightly to make the buyer's incentive strict, the equilibrium cutoff must:

$$w_k - p_k \in \delta[w_k - \bar{m}_S^k(y_S^k(w_k))],$$

equivalently,

$$p_k = (1 - \delta)w_k + \delta \bar{m}_S^k(y_S^k(w_k)), \quad (10)$$

where $\bar{m}_S^k(\cdot)$ is the convex hull of $m_S^k(\cdot)$, and

$$m_S^k(b) = \arg \max_{p_{k+1}} \left\{ p_{k+1} (G_S^k(b) - G_S^k(y_S^{k+1}(w_{k+1}(p_{k+1})))) + \delta V^{k+2}(G_S^{y_S^{k+1}, k}(\cdot)) \right\}.$$

The objective in m_S^k has strict single-crossing in (p_{k+1}, b) via the term $p_{k+1} G_S^k(b)$, so by the Monotone Selection Theorem (Milgrom and Shannon, 1994), any selection from $m_S^k(b)$ is non-decreasing.

Since $y_S^k(w_k)$ is continuous and strictly increasing in w_k , there exists a unique continuous, non-decreasing function $w_k(p_k)$ satisfying (10).³¹ Thus w_k is uniquely determined by p_k . To induce cutoff $w_k = c$, the seller posts $p_k = \max w_k^{-1}(c)$, which maximizes current profit among prices inducing the same current purchasing set. The standard no-randomization argument for seller-offer durable-goods bargaining then implies that the seller does not randomize on path after the initial period; see Ausubel and Deneckere (1989, Proposition 4.3).³² Thus, the indifference condition simplifies to:

$$w_k - p_k = \mathbb{E}_{\sigma_{k+1}}[\delta(w_k - p_{k+1})] \Rightarrow w_k - p_k = \delta(w_k - p_{k+1}),$$

consistent with Fudenberg et al. (1985) and Gul et al. (1986).

The inductive step for the seller. Consider the seller's pricing decision in period $t = k$ following the history p^{k-1} .

Case 1: Market should have cleared but has not. If, under the equilibrium strategies (Π, α) (allowing for possible seller deviation), the market is predicted to have cleared yet continues, this history is off-path. As before, we may assign arbitrary beliefs and allow the seller to post any price that deters such deviations.

Case 2: Market not cleared on-path. If the market is not cleared and this is consistent with the equilibrium path, the remaining posterior over v is F_S^k with unconditional CDF G_S^k . For any p_k in the support of σ_k we have:

$$p_k \in M_k(G_S^k) = \arg \max_{p_k} \left[p_k (G_S^k(\bar{v}) - G_S^k(y_S^k(w_k(p_k)))) + \delta V^{k+1}(G_S^{y_S^k, k+1}(\cdot)) \right],$$

and the value function satisfies:

$$V^k(G_S^k) = \max_{p_k} \left[p_k (G_S^k(\bar{v}) - G_S^k(y_S^k(w_k(p_k)))) + \delta V^{k+1}(G_S^{y_S^k, k+1}(\cdot)) \right].$$

By the induction hypothesis:

- $w_k(p_k)$ is continuous in p_k ,
- $y_S^k(w_k)$ is continuous in w_k ,
- $G_S^{y, k+1}(\cdot)$ is continuous in y , and
- $V^{k+1}(\cdot)$ is continuous.

³¹Existence follows from Berge's Maximum Theorem and upper-hemicontinuity.

³²Seller randomization after deviations may still be necessary, as in standard equilibrium constructions.

Therefore the objective

$$p_k (G_S^k(\bar{v}) - G_S^k(y_S^k(w_k(p_k)))) + \delta V^{k+1}(G_S^{y_S^k, k+1}(\cdot))$$

is jointly continuous in (p_k, G_S^k) . By Berge's Maximum Theorem, $V^k(G_S^k)$ is continuous in G_S^k , and $M_k(G_S^k)$ is non-empty and compact. Because the period- k information structure leaves the delaying buyer indifferent, her payoff equals that from buying immediately in period k at price p_k . Thus, at $t = k - 1$ the indifference condition is:

$$w_{k-1} - p_{k-1} = \delta \cdot \mathbb{E}_{\sigma_k}[w_{k-1} - p_k],$$

where w_{k-1} defined this way is the unique cutoff anticipating continuation play. Since delaying is optimal for the marginal buyer, her continuation payoff is nonnegative, so $w_{k-1} \geq \mathbb{E}_{\sigma_k}[p_k]$. Rearranging the indifference condition gives:

$$p_{k-1} = (1 - \delta)w_{k-1} + \delta \mathbb{E}_{\sigma_k}[p_k] \geq \mathbb{E}_{\sigma_k}[p_k].$$

Iterating the seller's inductive step together with the inductive step for Nature from $t = T$ backward completes the proof.

Step 4. Equilibrium strategy profile. The constructed equilibrium has the following on-path strategies:

- *Deterministic play on path.* The strategies of the buyer, the seller, and Nature are deterministic on the equilibrium path, following possible seller randomization in period $t = 1$.
- *Nature's strategy.* In each period t —whether h_S^t is on-path or off-path—Nature reveals to the buyer whether

$$v \leq y_S^t(w_t(p_t)).$$

- *Seller's strategy.* In period t , the seller chooses p_t to solve

$$\max_{p_t} \sum_{s=t}^T \delta^{s-t} p_s(p_t) \frac{F(y(w_{s-1}(p_t))) - F(y(w_s(p_t)))}{F(y(w_{t-1}))},$$

where $p_s(p_t)$ denotes the price posted in period s along the continuation path following p_t in period t .

- *Buyer's strategy.* In period t , the buyer purchases if and only if

$$v > y_t,$$

given the signal from Nature; otherwise, she waits. $\square$

Proof of Theorem 2. The proof proceeds in three steps:

Step 1. Fixing a price path and admissible threshold processes. Fix an arbitrary deterministic price path $(p_1, p_2, \dots)$. By Proposition 3 in Libgober and Mu (2021), the worst-case information structure against any fixed price path can be taken to be a *threshold* process (not necessarily myopic). Hence Nature's choice is summarized by a nonincreasing sequence $(y_t)_{t \geq 1}$, with the buyer purchasing at the first t such that $v > y_t$.

Under any such process, at time t the IC constraint is

$$\int_{\underline{v}}^{y_t} (v - p_t) f(v) dv \leq \sum_{s=t+1}^{\bar{T}} \delta^{s-t} \int_{y_s}^{y_{s-1}} (v - p_s) f(v) dv, \quad (11)$$

and the seller's discounted profit is

$$\sum_{s=1}^{\bar{T}} \delta^{s-1} \int_{y_s}^{y_{s-1}} p_s f(v) dv. \quad (12)$$

Our goal is to show that for every t either (i) $y_t = y_{t-1}$, or (ii) (11) binds at t .

Step 2. Block decomposition of the price path and Pooling thresholds within a block.

Let $p_0 \equiv +\infty$ and define recursively the block start times

$$t_1 := 1, \quad t_k := \min\{t > t_{k-1} : \delta^{t-t_{k-1}} p_t < p_{t_{k-1}}\}, \quad k \geq 2.$$

Thus, within each block $\{t_k, \dots, t_{k+1} - 1\}$ we have the monotonicity of *discounted* prices:

$$p_{t_k} \leq \delta p_{t_{k+1}} \leq \dots \leq \delta^{t_{k+1}-t_k-1} p_{t_{k+1}-1}. \quad (13)$$

We first prove that for any k , it is (weakly) optimal for Nature to *pool* the thresholds inside the block:

$$y_{t_k} = y_{t_{k+1}} = \dots = y_{t_{k+1}-1} \quad (14)$$

with the pooled level fixed at $y_{t_{k+1}-1}$.

IC under pooling. Start from any profile $\{y_{t_k}, y_{t_{k+1}}, \dots, y_{t_{k+1}-1}\}$ and replace it by (14) with the pooled level fixed at $y_{t_{k+1}-1}$. By (13), for every v and every $\ell \in \{1, \dots, t_{k+1} - t_k - 1\}$,

$$(v - p_{t_k}) \geq \delta^\ell (v - p_{t_k+\ell}).$$

with strict inequality whenever $v > 0$. Hence any type who was recommended to buy at some $t_k + \ell$ inside the block ($y_{t_k} \geq v > y_{t_{k+1}-1}$) weakly prefers buying already at t_k after pooling, and strictly prefers doing so whenever $v > 0$; any type with $v \leq y_{t_{k+1}-1}$ still learns “wait” at t_k and prefers to wait. Thus recommendations of waiting remain obedient after pooling.

Profit under pooling. Pooling moves all within-block purchases forward to t_k . For any type that originally bought at $t_k + \ell$, the seller’s discounted revenue changes from $\delta^\ell p_{t_k+\ell}$ to p_{t_k} , which weakly *decreases* by (13), strictly if some inequality is strict. If the pooling (by raising the value of waiting) causes some types to *defer beyond the block*, profit weakly decreases further because purchases can occur only at block starts $\{t_1, t_2, \dots\}$, and these satisfy

$$p_{t_1} > \delta^{t_2-t_1} p_{t_2} > \delta^{t_3-t_1} p_{t_3} > \dots,$$

so shifting a purchase from t_k to a later block $t_{k'}$ lowers discounted revenue. Therefore pooling within each block is (weakly) profitable for Nature and (weakly) reduces the seller’s profit. Thus, it is without loss to assume there is no information (as the thresholds are equal) and no sale within each block.

After pooling, Nature’s problem reduces to choosing the block–start thresholds $\{y_{t_k}\}_{k \geq 1}$ subject to the (block–level) IC constraints

$$\int_{\underline{v}}^{y_{t_k}} (v - p_{t_k}) f(v) dv \leq \sum_{s=k+1}^{\infty} \delta^{t_s-t_k} \int_{y_{t_s}}^{y_{t_{s-1}}} (v - p_{t_s}) f(v) dv, \quad (15)$$

(where the sum is finite if the horizon is finite).

Of course, this is slightly more general than required for the proof of Theorem 2, since the equilibrium pricing path in Theorem 1 automatically satisfies $p_1 > \delta p_2 > \delta p_3 > \dots$.

Step 3. Identifying the binding constraints. Define $\bar{v}_{t_k}$ such that

$$(1 - \delta^{t_{k+1}-t_k}) \bar{v}_{t_k} := p_{t_k} - \delta^{t_{k+1}-t_k} p_{t_{k+1}}$$

and $\bar{y}_{t_k}$ such that

$$\mathbb{E}[v \mid v \leq \bar{y}_{t_k}] = \bar{v}_{t_k}.$$

In other words, $\bar{v}_{t_k}$ is the expected valuation indifferent between buying at t_k at price p_{t_k} , and buying at t_{k+1} at price $p_{t_{k+1}}$. Because in Step 1 we established that

$$p_{t_1} > \delta^{t_2-t_1} p_{t_2} > \delta^{t_3-t_1} p_{t_3} > \dots,$$

it follows that $\bar{v}_{t_k} > 0$. In particular, define $\bar{y}_T$ such that $\mathbb{E}[v \mid v \leq \bar{y}_T] = p_T$.

We aim to prove that, Nature's optimal threshold information arrival process takes the following form: for every t_k , either $y_{t_k} = y_{t_{k-1}}$ or $y_{t_k} \geq \bar{y}_{t_k}$. If this is the case, then at any t_k , there is no information value in the future, and we must have $y_{t_k} = \bar{y}_{t_k}$ as $y_{t_k} > \bar{y}_{t_k}$ will violate the IC constraint.

Perturbation idea. We construct a perturbation of (y_{t_j}, y_{t_k}) that strictly reduces the seller's profit if both of the following two conditions holds at t_k :

$$(a) \ y_{t_k} < \bar{y}_{t_k}, \quad (b) \ y_{t_k} < y_{t_{k-1}}.$$

We will use $t_k = T$ as an example, but the same argument works for all t_k . Suppose $t_k = T$, and assume all previous IC constraints are slack. Then increasing y_{t_k} reduces the seller's profit, since some mass of buyers who would have purchased at t_{k-1} delay to t_k , and we already know $p_{t_{k-1}} > \delta p_{t_k}$. The increase of y_{t_k} is always feasible while leaving all other thresholds y_t unchanged, until one of the following occurs: (a) $y_{t_k} = y_{t_{k-1}}$; (b) $y_{t_k} \geq \bar{y}_{t_k}$; (c) Some earlier IC constraint binds. If either of the first two cases occurs, the claim follows immediately. For the third case, note that because multiple earlier IC constraints could bind simultaneously, we focus on the binding IC with the largest time index, denoted by t_j .

At t_j , the binding IC condition is

$$\int_{\underline{v}}^{y_{t_j}} (v - p_{t_j}) f(v) dv = \sum_{s=j+1}^{\infty} \delta^{t_s - t_j} \int_{y_{t_s}}^{y_{t_{s-1}}} (v - p_{t_s}) f(v) dv. \quad (16)$$

We now increase y_{t_k} while simultaneously adjusting y_{t_j} so that (16) remains satisfied. Because the IC at t_j stays binding, all earlier IC constraints remain valid: from the buyer's perspective, there is no information value after t_j , since she is indifferent between buying and waiting at t_j . Moreover, all later IC constraints remain locally satisfied because t_j was chosen as the largest time index with a binding IC, implying that all subsequent IC are slack.

Derivative computations. Let $y_{t_j}(y_{t_k})$ denote the t_j -threshold that satisfies (16) given y_{t_k} . Differentiate (16) w.r.t. y_{t_k} , holding other y_{t_s} fixed for $s \neq j, k$:

From the y_{t_k} term on the RHS of (16):

$$\delta^{t_k - t_j} (-(y_{t_k} - p_{t_k})) f(y_{t_k}). \quad (17)$$

From the y_{t_j} term: Move the y_{t_j} -integral on the RHS of (17) to the LHS and differentiate:

$$(y_{t_j} - p_{t_j})f(y_{t_j}) - \delta^{t_{j+1}-t_j}(y_{t_j} - p_{t_{j+1}})f(y_{t_j}) = (1 - \delta^{t_{j+1}-t_j})(y_{t_j} - \bar{v}_{t_j})f(y_{t_j}), \quad (18)$$

By the chain rule, keeping (16) binding requires that (17) equals (18) multiplied by $y'_{t_j}(y_{t_k})$:

$$\delta^{t_k-t_j}(p_{t_k} - y_{t_k})f(y_{t_k}) = (1 - \delta^{t_{j+1}-t_j})(y_{t_j} - \bar{v}_{t_j})f(y_{t_j})y'_{t_j}(y_{t_k}). \quad (19)$$

We claim that $y_{t_j} > \bar{v}_{t_j}$. Indeed, by hypothesis the IC at t_j binds. Recall that $\bar{v}_{t_j}$ is defined as the expected value at which the buyer is indifferent between buying and waiting at t_j even if no future information arrives. If instead $y_{t_j} \leq \bar{v}_{t_j}$, then any buyer told $v \leq y_{t_j}$ would strictly prefer to wait at least until t_{j+1} , contradicting the binding of the IC at t_j .

This establishes the derivative relationship (19) with $(y_{t_j} - \bar{v}_{t_j}) > 0$, which is the key ingredient for the profit-sign calculation in the next step.

Effect of the perturbation on profit. We now differentiate the seller's profit (12) under the perturbation of (y_{t_j}, y_{t_k}) constructed above. Since all other thresholds are fixed, only the t_j -, t_{j+1} -, and t_k -terms in (12) vary. Thus it suffices to differentiate

$$p_{t_j}(1 - F(y_{t_j}(y_{t_k}))) + \delta^{t_{j+1}-t_j}p_{t_{j+1}}F(y_{t_j}(y_{t_k})) - \delta^{t_k-t_j}p_{t_k}F(y_{t_k})$$

where we have made explicit the dependence of y_{t_j} on y_{t_k} . Differentiating term-by-term with respect to y_{t_k} yields:

$$-p_{t_j}f(y_{t_j}(y_{t_k}))y'_{t_j}(y_{t_k}) + \delta^{t_{j+1}-t_j}p_{t_{j+1}}f(y_{t_j}(y_{t_k}))y'_{t_j}(y_{t_k}) - \delta^{t_k-t_j}p_{t_k}f(y_{t_k})$$

Multiply through by $(y_{t_j} - \bar{v}_{t_j}) > 0$ (recall this sign was established above) and use (19) to substitute for $y'_{t_j}(y_{t_k})$ wherever it appears. After straightforward algebra and factoring out common positive terms, the derivative of profit with respect to y_{t_k} is proportional to

$$\left(-p_{t_j} + \delta^{t_{j+1}-t_j}p_{t_{j+1}}\right) \frac{\delta^{t_k-t_j}}{1 - \delta^{t_{j+1}-t_j}}(p_{t_k} - y_{t_k}) - \delta^{t_k-t_j}p_{t_k}(y_{t_j} - \bar{v}_{t_j}).$$

Divide by $\delta^{t_k-t_j}$ and substitute the definitions of $\bar{v}_{t_k}$. The change in profit from increasing y_{t_k} is proportional to

$$\bar{v}_{t_j}(y_{t_k} - p_{t_k}) + p_{t_k}(\bar{v}_{t_j} - y_{t_j}) = \bar{v}_{t_j}y_{t_k} - p_{t_k}y_{t_j}. \quad (20)$$

Case analysis.

Case 1: $p_{t_k} > \bar{v}_{t_j}$. Since $y_{t_j} \geq y_{t_k}$, (20) is strictly negative. Profit is reduced by increasing y_{t_k} .

This is also slightly more general than required for the proof of Theorem 2, since the equilibrium pricing path in Theorem 1 automatically satisfies $p_{t_k} < \bar{v}_{t_j}$.

Case 2: $p_{t_k} \leq \bar{v}_{t_j}$. We want to prove $\bar{v}_{t_j} y_{t_k} - p_{t_k} y_{t_j}$ is globally negative. Threshold-ratio monotonicity says $v \mapsto v/y(v)$ is decreasing, so $\frac{p_{t_k}}{y_{t_k}} \geq \frac{\bar{v}_{t_j}}{y_{t_j}}$, which implies

$$\bar{v}_{t_j} \bar{y}_{t_k} - p_{t_k} \bar{y}_{t_j} \leq 0.$$

Note we must have $y_{t_j} \geq \bar{y}_{t_j}$. Since the IC at t_j is binding, if $y_{t_j} < \bar{y}_{t_j}$, then any buyer recommended to wait would strictly prefer waiting even in the absence of future information—a contradiction.

Thus, we conclude globally that (given $y_{t_k} < \bar{y}_{t_k}$)

$$\bar{v}_{t_j} y_{t_k} - p_{t_k} y_{t_j} < 0,$$

and it is optimal for Nature to increase y_{t_k} while adjusting y_{t_j} until one of the following occurs: (a) $y_{t_k} = y_{t_{k-1}}$; (b) $y_{t_k} \geq \bar{y}_{t_k}$; (c) $y_{t_j} = y_{t_{j+1}}$; (d) Some constraint between t_j and t_k binds. If either of the first two cases occurs, the claim follows. The third case is simply a special instance of the fourth: it means we can identify a larger time index and repeat the same procedure. Since the game has only finitely many periods, and the largest binding constraint can only shift forward in time, eventually we must reach a point at which either of the first two cases occurs.

For $t_k \neq T$, the argument is essentially the same as before except at the last period T , adjusting y_T has no effect on future profits (as T is the last period), while in earlier periods we must adjust for the impact on future profits. The logic, however, is unchanged.

By similar algebra, the change in profit from increasing y_{t_k} is proportional to

$$\bar{v}_{t_j} y_{t_k} - \bar{v}_{t_k} y_{t_j}.$$

By threshold-ratio monotonicity, the map $v \mapsto v/y(v)$ is decreasing, so

$$\frac{\bar{v}_{t_k}}{y_{t_k}} \geq \frac{\bar{v}_{t_j}}{y_{t_j}},$$

which implies $\bar{v}_{t_j} \bar{y}_{t_k} - \bar{v}_{t_k} \bar{y}_{t_j} \leq 0$. Similarly, we must have $y_{t_j} \geq \bar{y}_{t_j}$. Since the IC at t_j is binding, if $y_{t_j} < \bar{y}_{t_j}$, then any buyer recommended to wait would strictly prefer waiting even without future information—a contradiction. Thus, we conclude globally that (given $y_{t_k} < \bar{y}_{t_k}$)

$$\bar{v}_{t_j} y_{t_k} - \bar{v}_{t_k} y_{t_j} < 0,$$

and it is optimal for Nature to increase y_{t_k} while adjusting y_{t_j} . □

B. ADDITIONAL PROOFS

B.1. Payoff Guarantee Interpretation

Proof of Proposition 1. We prove the result by backward induction. Let the Nash equilibrium strategy be denoted by (σ^*, Π^*) . At time $t = T$, since Nature can secretly deviate, for any p_T in the support of $\sigma^*(h_S^T)$ the information structure $\Pi^*(h_N^T)$ must minimize the seller's expected payoff. For off-path p_T , it is also without loss of generality to assume that $\Pi^*(h_N^T)$ minimizes the seller's expected payoff, as this does not affect the equilibrium outcome. Given Nature's strategy Π^* , since the seller can also secretly deviate at $t = T$, $\sigma^*(h_S^T)$ must maximize the seller's expected payoff given Π^* . Repeating this reasoning recursively, the conclusion follows by backward induction. $\square$

Proof of Proposition 4. In any Nash equilibrium, the information arrival process must minimize the seller's ex-ante profit; otherwise Nature could deviate to another information arrival process that lowers the seller's payoff with this deviation observed by the buyer. Now consider any sequentially worst-case information arrival process Π^* supporting a pricing strategy σ^* that is a safe solution. For the fixed process Π^* , the buyer's equilibrium strategy remains optimal even if it can depend on the entire information arrival process when Nature's choice is Π^* . Hence Π^* yields the original equilibrium payoff in the hidden-price-observed-information normal form; by safety, this payoff is the one-shot worst-case payoff, so Π^* is itself a one-shot worst-case information arrival process. Moreover, the seller and buyer are sequentially rational at every history. Since future prices remain hidden, the same hidden-price argument as in Proposition 1 implies that the seller's complete pricing strategy is a best response to Π^* . But because Π^* is a complete contingent strategy, holding Nature's strategy fixed still allows its prescribed information structure to vary with the history generated by a seller deviation. Since the choice of σ^* is optimal against Π^* , we therefore obtain a Nash equilibrium of the hidden-price-observed-information normal form game. $\square$

B.2. Other Results on Equilibrium Characterizations

Proof of Proposition 2.

Step 1. Threshold information. We first show that in any monotone equilibrium, Nature must employ a threshold information structure in every period.

Fix an arbitrary monotone equilibrium and define

$$\lambda_t := \Pr [\text{buyer is recommended to purchase in period } t \mid p^t],$$

the on-path probability—conditional on the realized price history p^t —that Nature recommends purchase in period t . Let λ_∞ be the probability that the buyer is never recommended to purchase, and let y_∞ denote the buyer’s expected valuation conditional on this event. We can construct a sequence of price-dependent thresholds

$$\infty = v_0 \geq v_1 \geq v_2 \geq \cdots \geq v_\infty = 0$$

such that each v_t depends only on the realized price path p^t , and satisfies

$$\Pr[v_t < v \leq v_{t-1} \mid p^t] = \lambda_t$$

for every t and every price history p^t . Under this process, in period t Nature recommends “buy” if $v > v_t$ and “wait” if $v \leq v_t$.

Consider a deviation where Nature uses this threshold-based rule in every period. Since the seller observes only the price path, he cannot detect this deviation, so the price sequence is unchanged. For any history p^t , a buyer recommended to wait then faces a posterior distribution which is (by construction) *inferior in the FOSD sense* to the posterior distribution under the original equilibrium. By the monotonicity property, such buyers will still obey the “wait” recommendation. Under the deviation, for every period t and price history p^t , at least as much buyer mass is deferred as in the original equilibrium. Hence total social surplus weakly decreases. Buyers who are recommended to purchase in a given period retain the option to buy immediately or wait, so their utility does not decline. Since buyers’ utilities do not fall and total surplus weakly decreases, the seller’s payoff must fall—implying that Nature’s payoff increases. Therefore, in any monotone equilibrium, Nature’s on-path strategy can be taken to be a threshold rule in every period.

Step 2. Market clearing in finite time under threshold information. Using the notation from above, for any b let F^b denote the cdf of the lower part of the prior F truncated at b . We first show there exists b^* such that if the seller faces posterior F^{b^*} , he optimally posts price $\underline{v}$ and clears the market. Suppose the seller faces F^b at h_S^t . If he charges p_t , one feasible strategy for Nature is to *fully reveal* the buyer’s type. Then:

$$V(F^b) \leq p_t(F(b) - F(w_t)) + \delta F(w_t)w_t \leq w_t(F(b) - F(w_t)) + \delta F(w_t)w_t,$$

where the second inequality uses $p_t \leq w_t$. On the other hand, by posting $\underline{v}$ and selling to all remaining buyers,

$$V(F^b) \geq F(b)\underline{v}.$$

Note that $V(\cdot)$ is defined relative to a particular equilibrium (since multiple equilibria may exist).

Because F^{-1} is Lipschitz-continuous at 0, there exists $q^* > 0$ and $L < \infty$ such that

$$F^{-1}(q) - \underline{v} \leq Lq \quad \forall q \in [0, q^*],$$

which implies

$$v - \underline{v} \leq L F(v) \quad \forall v \in [\underline{v}, F^{-1}(q^*)].$$

Take $b \leq F^{-1}(q^*)$. Combining the above bounds on $V(F^b)$:

$$\begin{aligned} 0 &\geq F(b) \underline{v} - w_t (F(b) - F(w_t)) - \delta F(w_t) w_t \\ &\geq F(b) \underline{v} - (LF(w_t) + \underline{v}) (F(b) - F(w_t)) - \delta F(w_t) (LF(w_t) + \underline{v}) \\ &\geq (1 - \delta) F(w_t) \underline{v} - LF(w_t) (F(b) - F(w_t)) - \delta LF(w_t)^2 \\ &\geq F(w_t) [(1 - \delta) \underline{v} - LF(b) + LF(w_t) - \delta LF(w_t)] \\ &\geq F(w_t) [(1 - \delta) \underline{v} - LF(b) - \delta LF(b)]. \end{aligned}$$

The final term is positive for b sufficiently small, which forces $w_t = \underline{v}$ and hence $p_t = \underline{v}$. Thus there exists b^* such that if the seller's posterior is F^{b^*} , he clears the market by setting $p_t = \underline{v}$.

From thresholds to finite-time clearing. Suppose Nature uses a threshold information arrival process on-path with thresholds

$$\infty = v_0 \geq v_1 \geq v_2 \geq \dots \geq v_\infty = 0.$$

By Nature's sequential rationality, in each t buyers recommended to wait do so. Let y_t be the upper bound of the support of the posterior at time t . Starting from any $y > b^*$, and for any $\epsilon > 0$, there exists finite k such that ϵ mass of buyers exit the market within k periods. If not, then:

$$V(F^y) \leq \epsilon \bar{v} + \delta^k \bar{v},$$

which can be made arbitrarily small as $\epsilon \rightarrow 0$ and $k \rightarrow \infty$. But also

$$V(F^y) \geq F(y) \underline{v} \geq F(b^*) \underline{v} > 0,$$

a contradiction. Therefore, the market must be cleared within

$$\bar{T}(\delta) = \left\lceil \frac{k(1 - F(b^*))}{\epsilon} + 1 \right\rceil$$

periods.

Step 3. Backward induction. Since the market clears in finite time after any history, we can solve for the equilibrium by backward induction on t and the finite T exactly as in the proof of Theorem 1. This induction terminates at $T = \bar{T}(\delta)$. □

Proof of Corollary 1. Note that for any given deterministic price path, Theorem 2 and threshold-ratio monotonicity imply that the worst-case information arrival process is the myopic threshold process. We claim that Nature can implement this process *sequentially*, without committing ex ante: in each period t_k , by construction

$$\mathbb{E}[v \mid v \leq y_{t_k}] = \bar{v}_{t_k},$$

and the buyer is weakly indifferent between purchasing at t_k and waiting until t_{k+1} even if she conjectures that no further information will arrive.

Thus, for any deterministic price path generated by an equilibrium strategy, Nature would deviate to the myopic threshold process, and hence in any deterministic equilibrium Nature uses the myopic threshold process on path. Under this condition, solving for equilibrium reduces to determining the seller's optimal deterministic price path subject to sequential rationality. The remaining steps follow exactly as in the previous proof. □

Proof of Proposition 6. When Nature implements myopic threshold information arrival process, the seller's discounted expected profit from time t onward (taking $w_0 = y^*(w_0) = \bar{v}$) is:

$$\sum_{s=t}^T \delta^{s-t} p_s \frac{F(y^*(w_{s-1})) - F(y^*(w_s))}{F(y^*(w_{t-1}))}.$$

From Gul et al. (1986) and Ausubel and Deneckere (1989), we know there exists a weak-Markov equilibrium which pins down the seller's equilibrium price path. Nature and the buyer have no incentive to deviate, by the same argument as in the proof of Theorem 1. □

Proof of Proposition 7. The argument is exactly the same as in the proof of Corollary 1. □

B.3. Distributional Assumptions Yielding Threshold-Ratio Monotonicity

Proof of Proposition 3. Our goal is to show that $\frac{w}{y(w)}$ is decreasing in w , or decreasing in $y(w)$ since it increases in w . Let $y = y(w)$, then $w = \mathbb{E}[v \mid v \leq y] = \frac{\int_{v \leq y} v f(v) dv}{F(y)}$ so that

$$\frac{w}{y} = \frac{\int_{v \leq y} v f(v) dv}{y F(y)}.$$

The derivative with respect to y is

$$\frac{\partial(w/y)}{\partial y} = \frac{yf(y) \cdot yF(y) - (yf(y) + F(y)) \cdot (\int_{v \leq y} vf(v) dv)}{y^2 F(y)^2}.$$

Rearranging, this derivative is non-positive if and only if

$$\int_{v \leq y} vf(v) dv \geq \frac{y^2 f(y) F(y)}{yf(y) + F(y)}.$$

The above inequality holds at $y = \underline{v}$, so a sufficient condition for it to hold at every y is that the derivatives of two sides are ordered. That is, we want

$$yf(y) \geq \left(\frac{y^2 f(y) F(y)}{yf(y) + F(y)} \right)'.$$

We can compute the derivative of $\frac{y^2 f(y) F(y)}{yf(y) + F(y)}$ to be

$$\frac{(yf(y) + F(y)) \cdot (2yf(y)F(y) + y^2 f'(y)F(y) + y^2 f(y)^2) - y^2 f(y)F(y) \cdot (2f(y) + yf'(y))}{(yf(y) + F(y))^2},$$

which simplifies to

$$\frac{y^3 f(y)^3 + y^2 f(y)^2 F(y) + y^2 f'(y) F(y)^2 + 2yf(y)F(y)^2}{(yf(y) + F(y))^2}.$$

This expression is smaller than $yf(y)$ if and only if

$$yf(y)(yf(y) + F(y))^2 \geq y^3 f(y)^3 + y^2 f(y)^2 F(y) + y^2 f'(y) F(y)^2 + 2yf(y)F(y)^2.$$

After some more algebra, the desired inequality becomes

$$y^2 f(y)^2 F(y) \geq y^2 f'(y) F(y)^2 + yf(y)F(y)^2.$$

Dividing both sides by $yF(y)$, this is equivalent to

$$yf(y)^2 \geq yf'(y)F(y) + f(y)F(y).$$

We can further divide both sides by $F(y)^2$ to arrive at

$$y \frac{f(y)^2}{F(y)^2} \geq y \frac{f'(y)}{F(y)} + \frac{f(y)}{F(y)}.$$

Let $h(y) = \frac{f(y)}{F(y)}$ with $h'(y) = \frac{f'(y)}{F(y)} - \frac{f(y)^2}{F(y)^2}$. The above inequality then becomes

$$yh'(y) + h(y) \leq 0.$$

Note that $yh'(y) + h(y)$ is the derivative of $yh(y)$, so this reduces to $yh(y)$ decreasing in y . $\square$

B.4. Constant Price Paths and Equilibrium Multiplicity from Exogenous Information Arrival Process

Proof of Proposition 5. We consider two cases for this proof: first the case where $T = \infty$, and then the modified argument for $T < \infty$. In both cases, we construct the following equilibrium:

- On-path, the seller posts a price equal to the buyer's expected value $\mathbb{E}_F[v]$, and no information is revealed.
- The buyer randomizes purchase with a probability to be specified—specifically, chosen so that the seller has incentives to follow the equilibrium strategy.
- If the seller deviates, the equilibrium reverts to the sequentially worst-case outcome described in Theorem 1/Proposition 2.

We now prove this profile forms an equilibrium. It is immediate that following a deviation by the seller, the future play constitutes an equilibrium, by Theorem 1/Proposition 2. The same holds on-path: since the buyer's purchasing decision does not depend on v , the on-path distribution of v conditional on not having purchased at time t is simply F . Thus, the buyer is indifferent between purchasing and delaying, as both deliver payoff 0, making them willing to randomize. Moreover, because we assume the buyer (strictly) randomizes, no profitable deviation is available to them, as all actions occur with positive probability on-path.

It remains to show that the seller does not prefer to deviate on-path, for appropriately chosen randomization probabilities. Let r^* denote the profit obtained in the equilibrium of Proposition 2/Theorem 1. The seller obtains at most r^* following any deviation; in particular, since the buyer's posterior distribution on-path is always F , and the horizon is infinite, this property holds at every time. Suppose we seek an equilibrium where the seller's continuation value is v^* at every point

in time, with $v^* > r^*$. In this case, set the buyer's purchase probability to be ρ in every period, where ρ satisfies

$$v^* = \rho \mathbb{E}_F[v] + (1 - \rho) \delta v^* \quad \Rightarrow \quad \rho = \frac{v^*(1 - \delta)}{\mathbb{E}_F[v] - \delta v^*},$$

with $\rho \in (0, 1)$ whenever $v^* \in (r^*, \mathbb{E}_F[v])$.

Thus, by charging $\mathbb{E}_F[v]$, the seller obtains a higher payoff than from deviating. This verifies the conditions in the statement of the proposition: (i) the seller uses a constant price path; (ii) the profit obtained is any $v^* \in (r^*, \mathbb{E}_F[v])$; and (iii) the market does not clear in any finite time, since ρ is constant and hence the probability the buyer has not purchased at or before time K is $(1 - \rho)^K > 0$.

For the $T < \infty$ case, define $v_T = \mathbb{E}_F[v]$ so that the buyer buys with probability 1 in the last period. Given any v_{t+1} with $t < T$, define v_t and ρ_t by

$$v_t = \rho_t \mathbb{E}_F[v] + (1 - \rho_t) \delta v_{t+1}.$$

Let r_t denote the $T = t$ equilibrium payoff identified in Theorem 1 with prior F , clearly $r_t < \mathbb{E}_F[v]$ unless it is optimal for the seller to clear the market at price $\underline{v}$ at $t = 1$. Hence, the equilibrium can take the same form as above, provided the sequence $v_1, \rho_1, v_2, \rho_2, \dots, v_{t-1}, \rho_{t-1}$ (with $v_T = \mathbb{E}_F[v]$) is such that $v_t \geq r_t$ for all t . This can be done by carefully choosing ρ_t close to 1 so each v_t is close enough to $\mathbb{E}_F[v]$. In this case, the seller obtains a higher payoff under the constant price path than from deviating, and the buyer remains indifferent between purchasing at any time and thus is willing to follow the mixed strategy. $\square$

C. ADDITIONAL EXAMPLES

C.1. Lowering Seller Profit Further under Ex-Ante (Commonly Known) Information

Here we walk through an example illustrating how some information arrival processes can deliver lower seller surplus compared to the sequential worst-case. Suppose $T = 2$ and $F \sim U[0, 2]$. Consider the following information arrival process:

- In the first period, the seller charges some p_1^* on-path; following any p_1 , the buyer learns whether $v > \tilde{v}(p_1)$ and buys if and only if it is. We leave $\tilde{v}(p_1)$ as to-be-specified for now.
- In the second period, the seller charges price $\frac{\tilde{v}(p_1)}{8}$, the buyer receives no additional information, and the buyer purchases.
- If the seller deviates in the second period to a price $\hat{p}_2 \neq \frac{\tilde{v}(p_1)}{8}$, the buyer learns whether or not $v > 2\hat{p}_2$.

Since all remaining buyers in the second period have $v \leq \tilde{v}(p_1)$, the above construction ensures that the seller has no (strictly) profitable second-period deviation following any first period price. Indeed, in $t = 2$, on-path the seller obtains profit $\frac{\tilde{v}(p_1)}{8} \cdot \frac{\tilde{v}(p_1)}{2}$, where $\frac{\tilde{v}(p_1)}{8}$ is the price and $\frac{\tilde{v}(p_1)}{2}$ is the probability that $v \leq \tilde{v}(p_1)$. But as shown in the main text, the best alternative price p_2 for the seller is $p_2 = \frac{\tilde{v}(p_1)}{4}$, which delivers the same profit level $\frac{\tilde{v}(p_1)}{4} \cdot \frac{\tilde{v}(p_1)}{4}$ where $\frac{\tilde{v}(p_1)}{4}$ is both the price and the probability that $v \leq \frac{\tilde{v}(p_1)}{2}$, the threshold that Nature would use. Since $\mathbb{E}[v \mid v < \tilde{v}] = \tilde{v}/2$, if the buyer learns that $v < \tilde{v}$ in period 1 and does not buy, then the buyer obtains $\frac{3\tilde{v}}{8}$ in the second period. Since *every* buyer with $v \leq \tilde{v}$ faces the same information set (and in particular, chooses the same action), the value of $\tilde{v}$ such that the buyer is indifferent between buying at time 1 and delaying purchase to time 2 satisfies:

$$\frac{\tilde{v}}{2} - p_1 = \delta \frac{3\tilde{v}}{8} \Rightarrow \tilde{v} = \frac{8p_1}{4 - 3\delta}.$$

Suppose that Nature, in the first period, tells the buyer whether her value is above or below $\frac{8p_1}{4 - 3\delta}$. Given this information structure (as well as understanding that the seller will follow the equilibrium strategy), the buyer will delay if told her value is below the threshold and not if it is above the threshold. Since the probability the buyer's value is above the first period threshold is $1 - \frac{4p_1}{4 - 3\delta}$ (since $v \sim U[0, 2]$), the seller's profit can be written as:

$$p_1 \left(1 - \frac{4p_1}{4 - 3\delta} \right) + \delta \frac{4p_1}{4 - 3\delta} \frac{p_1}{4 - 3\delta}$$

Take first order condition:

$$1 - \frac{8p_1}{4 - 3\delta} + \frac{8p_1\delta}{(4 - 3\delta)^2} = 0 \Rightarrow p_1 = \frac{(4 - 3\delta)^2}{32(1 - \delta)}.$$

Profit at this price is:

$$\frac{(4 - 3\delta)^2}{32(1 - \delta)} \left(1 - \frac{4(4 - 3\delta)}{32(1 - \delta)} \right) + \delta \frac{4(4 - 3\delta)^2}{(32(1 - \delta))^2} = \frac{(4 - 3\delta)^2(32(1 - \delta) - 4(4 - 3\delta) + 4\delta)}{(32(1 - \delta))^2} = \boxed{\frac{(4 - 3\delta)^2}{64(1 - \delta)}}$$

We check that this solution does indeed involve interior solution so first order condition is sufficient.

Given p_1 , we have $\tilde{v} = 2$ if:

$$1 - \frac{(4 - 3\delta)^2}{32(1 - \delta)} = \delta \frac{3}{4} \Rightarrow \delta = 4/5.$$

So, if $\delta < 4/5$, this scheme involves profit exactly as above. If $\delta \geq 4/5$, all buyers delay to the

second period and no sale occurs in the first period, meaning the total profit is $\delta/4$. The seller's optimal first-period price given this information process is therefore $p_1 = \frac{(4-3\delta)^2}{32(1-\delta)}$. In either case, profit is lower than in the sequentially worst-case solution. We emphasize that this example does not contradict Theorem 2, even if the prior F is still the uniform distribution and therefore satisfies threshold-ratio monotonicity. The reason is that in this example the seller and buyer act *after* observing Nature's strategy.

C.2. Safety is not Vacuous

Consider the discrete distribution where $v = 1$ with probability $1/2$ and $v = 0$ with complementary probability. The concavification arguments of Kamenica and Gentzkow (2011) immediately imply that the worst case makes the buyer indifferent between purchasing and not whenever recommended to not purchase. Therefore, in the static problem, given a price $p < 1/2$, the information structure recommends purchase with probability r when $v = 1$, where r satisfies

$$p = \frac{(1-r)q}{(1-r)q + 1 - q} \quad \Rightarrow \quad r = \frac{q - p}{q(1 - p)},$$

where q is the prior that $v = 1$.

Although our model assumed a continuous value distribution, this is not essential to obtain Theorem 1 in the two-period case. In the second period, Nature induces expectation p_2 , which generates no additional option value, and in the first period, Nature induces expectation $w(p_1)$, the indifference value for a consumer facing price p_1 . Given $w(p_1)$, the second-period price maximizes

$$p_2 \left(\frac{w(p_1) - p_2}{w(p_1)(1 - p_2)} \right),$$

since $w(p_1)$ is also the probability that $v = 1$ in period 2. Maximizing this over p_2 yields

$$p_2 = 1 - \sqrt{1 - w(p_1)}.$$

Using this, we can solve for $w(p_1)$ from the indifference condition

$$w(p_1) - p_1 = \delta(w(p_1) - p_2).$$

Given a solution for $w(p_1)$ (assuming it is interior), p_1 is then chosen to maximize

$$\frac{1}{2} \cdot p_1 \left(\frac{1/2 - w(p_1)}{(1/2)(1 - w(p_1))} \right) + \left(\frac{1}{2} + \frac{1}{2} \left(1 - \frac{1/2 - w(p_1)}{(1/2)(1 - w(p_1))} \right) \right) \cdot w(p_1) \delta p_2(p_1) \left(\frac{w(p_1) - p_2(p_1)}{w(p_1)(1 - p_2(p_1))} \right).$$

This expression can be maximized numerically; doing so for $\delta = 2/3$ yields the following solution:

$$p_1 \approx 0.2609, \quad w(p_1) \approx 0.3700, \quad p_2 \approx 0.2072, \quad \text{Seller Payoff} \approx 0.0763.$$

For this price path, we verify that the resulting solution is not safe. Suppose instead that the seller charged prices p_1 and p_2 as above, but Nature used an information structure that perfectly revealed the value to the buyer in the second period. In this case, the buyer would optimally delay, since when $\delta = 2/3$:

$$\frac{1}{2} - 0.2609 < \frac{1}{2} \cdot \frac{2}{3}(1 - 0.2072).$$

On the other hand, the seller's payoff under this alternative—where the buyer purchases at time 2 whenever $v = 1$ —is $(2/3)(1/2)p_2 \approx 0.0691 < 0.0763$. Thus, the fully worst-case information arrival process is not the one previously identified.

While threshold-ratio monotonicity is only defined for continuous distributions, we note that it will be violated for continuous distributions approximating this discrete case—for instance, by taking n even and sufficiently large and considering

$$f(v) = (v - 1/2)^n(1 + n)2^n.$$

Intuitively, for moderate values of v —say, in the range $[1/4, 1/3]$ —and for n very large, the threshold $y^*(v)$ will be very close to 1 for all values in this range. As a result, over this interval, $y(v)$ increases only slightly as v increases, even with large changes in v . Hence, the ratio $\frac{v}{y(v)}$ will increase as well.