Coasian Dynamics under Informational Robustness

JONATHAN LIBGOBER* AND XIAOSHENG MU**

*University of Southern California
**Princeton University

November 27, 2021

Preliminary and incomplete; comments welcome

ABSTRACT. This paper studies durable good monopoly without commitment under an informationally robust objective. A seller cannot commit to future prices and does not know the information arrival process available to a representative buyer. We first consider the case where the seller chooses prices to maximize her profit guarantee against a *time-consistent worst-case information structure*. We show the solution to this model is payoff-equivalent to a particular known-values environment, immediately delivering a sharp characterization of the equilibrium price paths. Furthermore, for a large class of environments, arbitrary (possibly time-inconsistent) information arrival processes would not lower the seller's profit as long as these prices are chosen. We call a price path with this property a *self-confirming solution*. As certain versions of the informationally robust objective under limited commitment may very well involve time-inconsistency, we posit that the notion of a self-confirming solution may be useful for researchers seeking to tractably analyze these settings.

KEYWORDS. Durable goods monopoly, limited commitment, dynamic informational robustness, self-confirming equilibrium.

CONTACT. LIBGOBER@USC.EDU, XMU@PRINCETON.EDU. Acknowledgements to be added.

1. INTRODUCTION

This paper adds the possibility of buyer learning to the classic durable goods monopoly environment where the seller lacks commitment to future prices. In our model, a buyer does not know their willingness-to-pay for a seller's product, but learns about it over time according to some information arrival process. Meanwhile, the seller chooses prices over time, but only has the ability to commit to a price during the period in which it is posted. The Coase conjecture states that, in the absence of learning, the monopolist should be expected to obtain arbitrarily low levels of profit. Intuitively, the monopolist would be expected to cannibalize their future demand; as a result, the buyer would anticipate future price drops, creating an additional incentive to delay. The monopolist would respond to this possibility by lowering the price in the initial period; provided players are sufficiently patient, this process leads to an unravelling whereby the monopolist is left charging a price close to the lowest possible buyer value. In this paper we refer to the corresponding price changes that emerge in equilibrium as *Coasian dynamics*.

It is well known that the Coase conjecture relies upon particular economically meaningful restrictions on the seller's environment. Our starting point is the observation that the assumption that the buyer knows their value for the seller's product is one of them. If the buyer does not know their value for the seller's product, then the possibility emerges that what they learn about it will interact with the seller's actions (for instance, since whether information could possibly be decision relevant will depend on the prices). In this paper, we observe that, if the information arrival process for the buyer is known to the seller, then this feature has the potential to completely reverse the Coase conjecture. While the Coase conjecture obtains under known values, information arrival can complicate the forces which deliver Coasian dynamics, to the point where the main conclusions are undone. While this observation does rely upon a particular information arrival process as possible, it is not obvious which constraints on information arrival can be imposed without ruling out economically meaningful possibilities of learning.

Along these lines, our interest in this paper is determining whether the same conclusion holds when the seller does not know how the buyer learns about the product. To formalize this, we seek to adopt a worst-case objective for the seller, so that at every point in time the seller seeks to maximize their payoff given some benchmark for the worst-case information arrival process. Recent years have seen a flourishing of theoretical work using robust objectives in order to describe how mechanism designers may deal with limited understanding (or limited confidence in their understanding) of their relevant environment. These papers have shown that often, such concerns can lead to mechanisms and policies which are much simpler to interpret compared to those that seek to optimize against the precise details of the environment. Often, these approaches yield a dramatic gain of tractability to the point where certain unstudied (or understudied) applications can be formally analyzed. The *informationally* robust objective is a particular special case, where a designer (often a seller) maximizes payoffs against a worst-case information structure. The robust approach is well-suited to cases where the uncertainty relates to the informational environment; indeed, a large part of its popularity is due to the influence of the *Wilson Critique*, which posited that the strong epistemic assumptions made by mechanism design have severely limited is applicability.

A significant novelty of our use of the robust approach in this setting arises due to the assumption that the seller lacks commitment. Despite the growth of this field and many important contributions, the robust approach significantly lags the classical Bayesian approach in terms of how to model commitment issues while seeking robustness. In classical mechanism design, once the commitment assumption is relaxed, the principal and the agent engage in a game, since the agent chooses actions anticipating that the principal may offer a different option to the agent in the future. While these settings are often difficult to study fully without significant added structure[1] (typically motivated by details of an application), there nevertheless appears to be a consensus approach about how to tackle such problems (namely, using the PBE solution concept).

It is safe to say that a similar understanding does (yet) not exist in the robust mechanism design literature. As Carroll (2019) notes, "non-Bayesian models do face some extra hurdles.... trying to write dynamic models with non-Bayesian decision makers leads to well-known problems of dynamic inconsistency, except in special cases (e.g., Epstein and Schneider (2007)). This may be one reason why there has been relatively little work to date on robust mechanism design in dynamic settings." By contrast, Bergemann and Valimaki (2019) point out the importance of understanding how to move away from assumptions of Bayesian mechanism design in dynamic settings, writing that the literature on dynamic mechanism design has so for involved "... Bayesian solutions and relied on a shared and common prior of all participating players. Yet, this clearly is a strong assumption and a natural question would be to what extent weaker informational assumptions, and corresponding solution concepts, could provide new insights into the format of dynamic mechanisms."

This paper makes progress toward filling the gap outlined by Bergemann and Valimaki (2019) in the context of durable goods sales with limited commitment. Insofar as limited commitment mechanism design settings often require specialization to solve for optimal equilibrium policies, there are are several reasons this is a natural first place to study. First, it is perhaps the most thoroughly studied setting with limited commitment, and it is well-understood what forces drive

[1]As an example, Pavan (2017) writes: "'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." By contrast, our paper allows for such information dynamics, albeit using a robust approach.

the solution under the Bayesian benchmark. As one of our results shows a tight analogy between a particular environment without information arrival or robustness concerns, we can therefore immediately import the economic intuition from past work to provide insights our setting. Second, for durable goods pricing, we can make the comparison to the commitment case more precisely. If seller could instead commit to their strategy, our main model coincides with the one studied in our previous paper, Libgober and Mu (2021).

Our first contribution is to identify a way of specifying the robust objective under limited commitment whereby the worst-case is time-consistent; i.e., the worst-case the seller anticipates tomorrow will still be the worst-case when tomorrow arrives. This circumvents the hurdle identified by Carroll (2019) for our problem. Specifically, our first result considers a benchmark where the seller chooses a price, assuming that at each point in the future, the buyer's information structure will be chosen to minimize the seller's profit *from that period on.* In this approach, we take the idea of nature as a player seriously (although we are happy to view "nature" as more of an expositional device to explain why time-consistency is maintained in this benchmark). In this case, we show that the equilibrium outcome essentially coincides with the unique[2] equilibrium outcome in a particular *known-values* environment (i.e., where the buyers know their willingness-to-pay for the product). To see why, notice that when there is only a single period to sell, whether the seller has commitment or not is irrelevant. Our past work, Libgober and Mu (2021), showed that the outcome in the single period setting is identical to that of a particular environment with known-values, under a transformation of the value distribution (a process we referred to as *pressing*). Our work in this paper shows that, provided the information structure at every period is chosen to minimize the continuation profit at every point in time, the key properties of the solution are maintained, despite some added technical details due to limited commitment.

This benchmark assumes that the information structure is "reoptimized" at every point in time, just as the seller "reoptimizes" prices. This specification eliminates from consideration information structures which would hurt the seller's overall profit, if these would possibly be more favorable to the seller in the future. To what extent is this solution plausible as a "true-worst case," that is, as maximizing the seller's profit against *all* possible information arrival processes? Our second contribution is to show that the answer to the question depends on how the seller *believes* the information structure is determined. Perhaps surprisingly, despite allowing reoptimization in our main benchmark, we illustrate a sense in which the solution is a true worst-case, under some assumptions about the prior distribution of consumer values. As perhaps anticipated by Bergemann and Valimaki (2019), we make this point formally by studying a distinct solution concept. Specifically, we introduce what we refer to as a *self-confirming solution*, and show that it

[2]To avoid conceptual difficulties with equilibrium multiplicity, in this paper we focus exclusively on the "gap" case.

is often satisfied by the benchmark solution outlined above. In a self-confirming equilibrium (first defined by Fudenberg and Levine (1993)), player strategies are optimal given some conjecture of their environment, with an added requirement that this conjecture is never disproven during the play of a game. That is, while the behavior may require a favorable *and counterfactual* conjecture about an opponent's behavior, such conjectures can persist if they are *not inconsistent* with what would be observed during the play of the game. Similarly, in a self-confirming solution, while the seller may be misspecified about how nature selects the buyer's information arrival process, this misspecification cannot possibly lead them to obtaining a lower payoff. In other words, the requirement states that, *if* the seller *believed* nature did not have commitment, then the profit *would not be lower* if the seller were incorrect and nature *did*, in fact, have the power to commit.

The sufficient condition we identify for our baseline solution to be a self-confirming one is fairly weak and satisfied by many natural value distributions. Roughly speaking, the requirement is that there is not too much mass "toward the top" of the prior value distribution. In these cases, a worse information structure may withhold information from buyers, in order to induce additional delay. We also show that this requirement will always be satisfied in the gap-case toward the bottom of the distribution; a corollary of this is the observation that, in any environment where the purchasing threshold become concentrated around the minimum of the value distribution (e.g., if buyers are much less patient than sellers), then our baseline model's solution will *always* be a self-confirming one.

We believe this issue is provocative because it suggests that the dynamic inconsistency for robust limited commitment objectives issue may be less severe than originally thought. While we focus on a particular setting in order to obtain a sharp characterization of the optimum (in the tradition of the literature on optimal mechanisms under limited commitment, referenced by Pavan (2017)), we see no reason why this solution concept could not be used elsewhere. Notice that a correctly specified Bayesian decisionmaker would vacuously choose a self-confirming solution. So, a natural question is whether other interesting mechanism design settings also possess intuitive self-confirming solutions. Our hope is that by providing some clear intuition for our baseline model's solution, we can suggest a path forward in order to fill this important gap between Bayesian and robust approaches.

1.1. Relevant Literature

The early literature on robust mechanism design was motivated in large part by the desire to move away from strong reliance on common knowledge assumptions implicit in Bayesian mechanism design (Bergemann and Morris (2005), Chung and Ely (2007)). While these papers focused on the "known values" case (i.e., where participants are assumed to know their values), other papers that followed this literature considered the case of "unknown values" case where the designer also faced uncertainty about what participants knew about their own values (Bergemann et al. (2017), Du (2018), Brooks and Du (2021), Brooks and Du (2020), Libgober and Mu (2021)).[3] Ultimately, the economic motivation behind these papers strikes us as entirely orthogonal to the issue of whether the designer has commitment power—namely, we see no reason that a designer who faces uncertainty about the environment should necessarily be able to commit to their mechanism, given that often they cannot.

Actually, this failure may be more severe than it seems. If anything, there is a tension between the use of a maxmin setting on the one hand and a stark reliance upon the commitment assumption on the other. One criticism of the literature on robust mechanism design with commitment is that it is not clear how a seller would reliably obtain commitment power without also obtaining some greater certainty in the environment.[4] Despite this, as far as we are aware, the vast literature on robust mechanism design in general and the informationally robust objective in particular has only focused on the cases where the designer has commitment power, and in fact mostly focused on cases where the environment is static by assumption.[5] These papers typically acknowledge that the maxmin assumption is quite strong, but view it as an important step to understand the implications of other strong assumptions behind the Bayesian framework (which, we should stress, we also view as convincing). As suggested above, however, the main difficulty in this approach relates to dynamic consistency issues with the maxmin objective, as highlighted by Epstein and Schneider (2007), among others.

Therefore, our paper is part of an agenda that seeks to resolve conceptual issues that arise when extending the robust framework to domains that have been productively analyzed under Bayesian objectives. Though to the best of our knowledge we are the first to study relaxing commitment, other work fits into this larger agenda as well. Bolte and Carroll (2020) study the problem of a principal who can choose investment in the course of interacting with an agent, and show this provides a foundation for linear contracts, echoing an earlier result of Carroll (2015). Ocampo Diaz and Marku (2019) also extend Carroll (2015), but this time to consider the case of

[3]See also Lopomo et al. (2020) for a generalization of the robust framework to accommodate more "intermediate" cases where the designer may not resolve uncertainty in the precise way as outlined in these other papers.

[4]Commitment is often justified using a "repeated interactions" microfoundation; namely, if the designer broke a commitment promise, then future participants could "punish" them by reverting to a non-commitment equilibrium where presumably the designer would do worse. Yet in this case, presumably the designer would also have a better understanding of their environment having engaged in it for a long period of time, and therefore may very well *wish* to reoptimize using their better knowledge of the environment.

[5]Interestingly, while not explicitly about the robust objective, perhaps closest to this particular problem is Ravid et al. (2020). They consider the case of buyer optimal information when the choice of information structure is not observed by the seller. If information choice is observed, the buyer optimal information structure is worst-case for the seller (Du (2018)). Relaxing the assumption that the buyer can commit to their information structure, as in Ravid et al. (2020), is similar to relaxing the assumption that nature can commit to choices under the robust objective.

competing principals in a common agency game. Both of these papers address a similar conceptual issue, namely how the strategic choices of the designer should interact with their corresponding use of the maxmin objective. However, in both of these papers, the "worst-case" is only considered once, and hence the issue of time inconsistency (which is the core of our exercise, and apparently a major roadblock to more work on the robust objective with limited commitment) does not arise.

A less related literature has considered mechanism design problems where *agents* (instead of the designers) have non-Bayesian preferences, including the maxmin case. The motivation of these papers is quite different from the robust mechanism design literature, however, since the concern there is how the designer should react to the presence of non-Bayesian buyers (Wolitzky (2016), Bose and Renou (2014), Di Tillio et al. (2017)). Some papers in this literature explicitly consider time-inconsistency issues under dynamic formulations, and demonstrate how a designer may be able to exploit this particular feature (Bose et al. (2006), Bose and Daripa (2009)). Notice that in our case, for similar reasons as in our prior work (see Appendix F of Libgober and Mu (2021)), the "Bayesian agent" case is also worst-case for the seller, even if additional ambiguity could be added.

Lastly, we mention that recent work has considered the sensitivity of the Coase conjecture to the presence of information arrival (though as highlighted by Pavan (2017), this seems relatively unexplored in other limited commitment settings). The key conclusions from the literature on the Coase conjecture with known values are outlined in Ausubel et al. (2002). Under somewhat restrictive assumptions on either the type distribution or the learning process, Duraj (2020), Laiho and Salmi (2020) and Lomys (2018) consider how the conclusion of the Coase conjecture may be influenced by the presence of learning. The reason is that the direction and magnitude of selection pressures are known to be fundamental in its conclusion (see Tirole (2016)). In contrast, we show that under the robust objective, such forces are essentially absent—at least, under the natural benchmark we highlight here. We note that information arrival more generally can influence the direction and magnitude of selection, and thus certain conclusions from this literature may be sensitive to the presence of information arrival without the worst-case objective.

2. MODEL

A seller of a durable good interacts with a buyer in discrete time until some terminal date T, which for simplicity we take to be finite but arbitrary. The buyer can purchase the good at any time $t = 1, \ldots, T$. The buyer's value v is drawn from a continuous distribution F which the buyer and seller commonly know. However, the buyer does not know v and instead learns about it over time—our assumptions on how the buyer learns is described in Section 2.1. Our main focus in this paper is on the "gap" case, where the cost of producing the good is 0 for the seller, and the support of F is bounded away from 0.

Within a period t, the seller chooses a price $p_t \in \mathbb{R}_+$, after which the buyer decides whether to purchase or not. We assume the seller does not have commitment—while she is able to choose the price at which the buyer would purchase at time t, she cannot commit to prices offered in future periods.

2.1. Information Structures

In every period before deciding whether to purchase the object or not, the buyer receives information about her value for it. Specifically, at time t, prior to the deciding whether to purchase the object, the buyer observes a signal $s_t \in S_t$ which is drawn according to an information structure $I_t(s_{t-1}) : V \to \Delta(S_t)$. We emphasize that the signal at time t can depend upon the signal history up until time t. Throughout the paper, we assume that I_t is observed by the buyer. We let $\mathcal{I}$ denote the space of all possible (static) information structures, and we let $\mathcal{I}^t$ denote the space of sequences of information structures between time 1 and time t.

2.2. Strategies, Payoffs and Equilibrium

A major conceptual difficulty with our exercise is that we seek to use the robust objective when the seller does not have commitment. Such a formulation is known to be somewhat elusive, as the worst information structure for the seller at some future time may differ from the worst case when that time arrives (i.e., the worst-case may be dynamically inconsistent under maxmin preferences).

To address this, we consider the case where we treat the information structure as chosen by an adversarial nature. Such an interpretation is common from the robust mechanism design literature; in our case, it is useful in that it forces the seller to have a dynamically consistent view of the buyer's information. Specifically, we posit that the information structure is determined according to the following game:

- Within each period, the seller first chooses the price to be charged in that period, p_t.
- Next, an adversarial nature chooses an information structure I_t for that period (so that information in a given period may depend on the price in that period)
- The buyer decides whether to purchase in that period, given the signal observed and the equilibrium strategies being used by the seller and nature. Let $\sigma_{s^{t-1},I^t} : S_t \to \Delta\{0,1\}$ denote the buyer's strategy (i.e., a probability of purchase as a function of time t signal) if the information structure sequence up until time t has been $I^t \in \mathcal{I}^t$ and the signal sequence prior to time t has been $s^{t-1} \in S^{t-1}$.

Payoffs in period t are discounted by a factor of δ^{t-s} relative to payoffs in period s. In any period where the buyer does not buy, the payoffs are 0. If the buyer does buy in period t at price p_t, then the time 1 utility obtained by the seller is $\delta^{t-1}p_t$ (and nature's utility is therefore $-\delta^t p_t$). The buyer's payoff from purchasing in period t is $\mathbb{E}[v - p_t \mid I_1, s_1, \ldots, I_t, s_t]$.

To summarize, the buyer's strategy must be such that if $\sigma_{s^{t-1},I^t}(s_t) = 1$, then:

$$\mathbb{E}[v - p_t \mid I_1, s_1, \ldots, I_t, s_t] \geq \mathbb{E}\left[\max_{\tau: t<\tau\leq T} \delta^\tau \mathbb{E}[v - p_\tau \mid I_1, s_1, \ldots, I_\tau, s_\tau]\right].$$

whereas if $\sigma_{s^{t-1},I^t}(s_t) = 0$ then this inequality is flipped. If the buyer purchases at some time s at a price of p_s, then from the perspective of time $t < s$ the seller obtains payoff $\delta^{s-t}p_s$. The seller therefore chooses the time t price, as a function of $p_1, \ldots, p_{t-1}$ alone, to maximize:

$$p_t \mathbb{P}[\sigma_{s^{t-1},I^t}(s_t) = 1 \mid p_1, \ldots, p_t, \sigma] + \sum_{k=t+1}^{T} \delta^{s-t+1} \mathbb{E}_{p_k \sim \sigma_s}[p_k \mathbb{P}[\sigma_{s^{k-1},I^k}(s_k) = 1 \mid p_1, \ldots, p_k, \sigma] \quad (1)$$

By contrast, at each time t, nature chooses the information structure to minimize the seller's payoff, given by (1).

2.3. Discussion

Our explicit use of "nature" as a player is primarily expositional device to explain why one might expect time consistency of the information structure to be maintained. In subgame perfect equilibrium, actions are required to maximize payoffs, given that future actions are determined according to the equilibrium profile (and in turn, these actions must satisfy the same requirements). Thus, when a player chooses an action, they do so (correctly) anticipating future actions, and do not change their conjecture of future actions when the future arrives. By framing the information structure choice as emerging from a game, we seek to highlight that this consistency is maintained. By contrast, in single-agent problems under a maxmin objective, generally it is possible for some action to be taken anticipating a future worst-case which would not actually be worst-case when the future arrives. This possibility we rule out in our model, though we will discuss the role of it again in our analysis.

A natural question is the solution to this model if the seller could commit to their choice of strategy at time 1, instead of having to reoptimize their choice at every time given the history of actions. This problem was solved in our prior paper, Libgober and Mu (2021), where we showed that the optimal selling strategy is a constant price path.[6] More precisely, we identified a *known-*

[6]If the seller commits to a pricing strategy, then the issue of nature commitment play any role.

values environment which would deliver an identical pricing strategy as optimal. Despite the similarity, the conceptual difficulties were essentially orthogonal to those studied here. For one, a major technical concern in that paper was the fact that with commitment, one might worry about the ability for the seller to randomize, even though such randomizations would never be necessary in a known-values model. Here, the fact that nature picks the information structure every period immediately shuts down any gains to randomization in any future period. On the other hand, in that paper the determination of the information structure was significantly more straightforward, in part since nature makes one choice and in part because the worst-case information structure is dramatically simpler against constant price paths. On the other hand, the distinction between the exercises disappears in the special case where $T = 1$, and for this reason we refer the reader to that paper for more details about the solution to that benchmark.

3. SOLUTION TO THE BASELINE MODEL

We now proceed to solve the previous model.[7] Note that in the $T = 1$ case, the issue of non-commitment does not arise, and the solution is exactly as articulated in Libgober and Mu (2021). Intuitively, one can use results from Bayesian persuasion to show that the worst-case information structure takes a partitional form, where the partition depends on the price charged by the seller. Using the mapping between prices and thresholds, one can then derive a value distribution which, under an assumption of *known* values, gives an identical solution to the seller's problem. We review the definition of this corresponding value distribution, which we dubbed the pressed-distribution:

Definition 1 (From Libgober and Mu (2021))**.** *Given a continuous distribution F, its "pressed version" G is another distribution defined as follows. For $y > \underline{v}$, let $L(y) = \mathbb{E}[v \mid v \leq y]$ denote the expected value (under F) conditional on the value not exceeding y. Then $G(\cdot) = F(L^{-1}(\cdot))$ is the distribution of $L(y)$ when y is drawn according to F.*

Now, in that paper, we also showed by example that one should generally not expect the pressed distribution to completely describe the seller's problem if they were to use a declining price path. The reason is that some information structures may lower the seller's profit by revealing more information to the buyer. Thus, in dynamic environments, it is not immediately clear that one can say that the seller's problem is "as-if known values under the pressed distribution." While that paper does feature constant price paths as delivering the optimum, this feature should decidedly not be the case given that we are focused on the noncommitment case.

[7]We briefly mention that the same results apply in the no-gap case with a finite horizon, though as is well-known under known values, the finite horizon assumption is more restrictive in the no-gap case than the gap case.

Our first result shows that those information structures are time-inconsistent, in that they rely upon giving the buyer more information than the worst-case at later times. If one forces those information structures to also minimize the seller's profit from that time on, then we again recover the tight analogy:

Theorem 1. *Equilibrium payoffs in the baseline game are unique. Furthermore, an equilibrium is given by the following:*

- *The information structure is partitional.*
- *The prices the seller charges coincide with the prices charged when the buyer's value is drawn according to the pressed version of F, and where the buyer knows his value.*

The intuition behind this result is fairly straightforward. At every time, nature chooses information to minimize the seller's total discounted payoff from that time on. Given this, in adjusting the threshold, nature knows that the next period choice of threshold depends only on the price the seller is expected to charge in that period. As a result, a small change in the threshold today would have no change in the threshold in the future, meaning that the optimal choice is simply to minimize the seller's expected profit from that period on. Note that a technical issue is that there may be multiple equilibria, as different information structure choices of nature might induce identical behavior from the buyer, as a function of the buyer's true value. However, we show that this possibility does not change the conclusion of the result.

The key property of this result is that the worst-case is time consistent. In the last period, say period T, the worst-case information structure involves a price-dependent threshold. In the next-to-last period, the equilibrium determines what the last period price should be. The seller anticipates that the worst-case information will be of a threshold form, with the threshold depending on this (anticipated) price. Crucially, the worst-case for I_T is both the worst case when period T begins, as well as at any $t < T$. This same reasoning applies to earlier information structures as well, although the thresholds for these information structures will depend on the value at which the buyer would be indifferent between purchasing and not, instead of the price.

Theorem 1 provides a sharp characterization of the equilibrium payoffs. The reason the *outcome* is not unique is due to the possibility that nature provides some richer information structure to the buyers, which nevertheless induces the same behavior. However, the result allows us to provide some sharp descriptions of the outcome in the worst-case. For instance, a corollary is that provided the pressed distribution satisfies certain weak conditions in Ausubel et al. (2002), there are a finite number of periods after which the market clears. This need not hold for an *arbitrary* (non-worst case) information arrival process. The issue more generally is that information arrival *in principle* can generate a gap between the seller's "on-path" payoff and the "off-path" punishment payoff. The existence of such a gap drives, for instance, the folk theorem of Ausubel and Deneckere (1989). This contrasts with stationary equilibria, such as the one in Theorem 1, where even off-path the strategy only depends on the size of the remaining market. As an example, consider the following proposition, which stands in stark contrast to the equilibrium outcomes in the known-values model:

Proposition 1. *Fix F, δ and time horizon K. Suppose the equilibrium outcome under known values with distriution F does not involve the market clearing at time 1. Then there exists an information structure, optimal stopping time for the buyer and equilibrium price path for the seller such that the market does not clear before time K.*

This proposition states that it is impossible to bound the time at which the market will clear over the set of all information structures. By contrast, a key result from the known-values gap case is that such a uniform time at which the market clears can be found, under general conditions. We therefore view this proposition as a proof of concept, illustrating the difficulty of deriving analogies between the Coasian known-values settings and those with information arrival in full generality. This observation highlights our claim that the robust approach is appealing in that it maintains analogies to the known-values case, and that certain conclusions should not immediately be taken for granted when seeking to accommodate information arrival into the Coasian setting without this approach.

4. RICHER NATURE COMMITMENT

Theorem 1 provides a striking characterization of the solution to the baseline model—it coincides with a certain known-values environment, which was previously identified in the commitment version of the same model. We have therefore identified an environment with limited seller commitment where the value of commitment under an informationally robust objective can be determined from the value of commitment under known values.

A natural question this raises is whether this is in fact a "true-worst case." To be more precise, note that our game features a timing protocol whereby the seller moves first in each period, and nature then responds. It is possible that, were nature able to pick their strategy *before* playing the game (so that the need to best reply to the seller were eliminated), the seller could be forced to an even lower profit. Can dropping the incentive constraints of nature hurt the seller even more?

There is a special case where it cannot, which is when the solution to the previous model involves $p_2 = \underline{v}$; that is, where the seller clears the market at time 2. This is straightforward to show—in this case, nature's choice does not influence behavior at time 2, and so its problem is essentially static. In this case, the problem of nature is essentially a Bayesian Persuasion problem, and in the environment we study the worst-case is known to take a threshold form, where the threshold is chosen so that a buyer who does not purchase is indifferent between actions.

More generally, the answer turns out to depend on what we assume about the seller's view of nature. Suppose we were to assume that the seller *knew* nature had such commitment power, and therefore chose their strategy to best respond to this (committed to) information arrival process. The theorem below shows that there does exist an information arrival process which delivers a lower profit.

Theorem 2. *Suppose the equilibrium outcome in Theorem 1 does not involve purchase by time 2 with probability 1. Then the seller's profit is lower when nature can commit to an information structure.*

The following example illustrates:

Example 1. *Suppose $T = 2$ and $v \sim U[0, 2]$.*[8] *Note that this implies the pressed distribution is $U[0, 1]$. We can therefore compute (see the Appendix for details) that the equilibrium to the baseline model involves the following as the solution for prices p_1, p_2 and seller profit, say π, as:*

$$p_1 = \frac{(2-\delta)^2}{8-6\delta}, p_2 = \frac{(2-\delta)}{8-6\delta}, \pi = \frac{(2-\delta)^2}{4(4-3\delta)}.$$

Moving back to the nature's original problem, the information structure that nature chooses tells the buyer at time 2 whether v is above or below $2p_2$; since, at time 1, a buyer with value $2p_2$ would be indifferent between purchasing and not, the time 1 threshold informs the buyer whether or not the value is above or below $4p_2$.

We now exhibit the information structure which holds the seller down to a lower profit. Let $\pi^*(\tilde{v})$ denote the seller's profit as a function of the first period threshold $\tilde{v}$, above which consumers learn their true *value and purchase (i.e., $\tilde{v}$ is not the indifferent value, but the partition threshold). Consider the following second period outcome:*

- *In the second period, following any first period history, the seller charges price $\pi^*(\tilde{v})$, nature provides no information, and the buyer purchases.*
- *If the seller deviates in the second period, nature uses the worst-case partitional threshold.*

By construction, the seller has no (strictly) profitable second period deviation, no matter what the first period price is. Furthermore, note that, since $\pi^*(\tilde{v}) < \mathbb{E}[v \mid v < \tilde{v}]$, the buyer is willing to follow this strategy as well. The calculation of the resulting optimal first period price is now similar

[8] While this violates our assumption that the value distribution satisfies the gap case, this for illustrative purposes and the same conclusion would hold for $v \sim U[\varepsilon, 1]$.

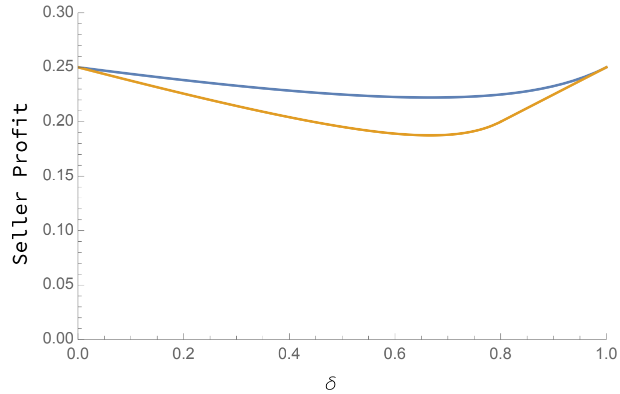

Figure 1: Comparison of seller profit between equilibrium of the baseline model vs. Theorem 2, for Example 1

to the previous case. The difference is in the calculation of the indifferent value in the first period, since the buyer now obtains additional surplus from delay. We can show that if nature were to choose an information structure of this *form, then the seller could prevent* all *sale in the first period when* $\delta \geq 4/5$ *(and in this case, the seller's expected profit is* $\delta/4$*, since the expected profit from the one period problem is* $1/4$*); otherwise, the seller's profit is:*

$$\frac{(4-3\delta)^2}{64(1-\delta)}.$$

Figure 1 plots, as a function of δ*, the profit the seller obtains in the equilibrium of the baseline model (blue line) to the profit the seller obtains in the equilibrium under this different information structure. We have that this is uniformly lower, except for when* $\delta = 0$ *and when* $\delta = 1$*, in which case the seller's problem is essentially static (with only the first period mattering in the former case and all sale happening in the second period in the latter case).*

The Theorem essentially generalizes the example to any setting where the market does not clear at time 2. The key point is that the solution to the baseline model leaves additional scope to transfer surplus to the buyer in order to induce additional delay. In the information structure nature chooses, the seller obtains the exact same continuation profit as in the baseline model, but the inefficiency entailed disappears. Instead, the buyer obtains more surplus, which makes them more willing to delay, thus hurting the seller's profit.

This result suggests that perhaps the solution to the previous model is not a "true worst-case." However, one criticism of the benchmark where nature has full commitment is that it requires extreme confidence from the seller regarding nature's choice of information structure. It seems reasonable to ask where this confidence would come from.

To analyze this question, we consider the following criterion on price paths:

Definition 2. *An optimal pricing strategy from the baseline model is a **self-confirming solution** if the seller's anticipated equilibrium profit is equal to the worst-case profit guarantee over the set of all dynamic information arrival process.*

We are not aware of this concept being studied elsewhere in the robust mechanism design literature. To maintain focus, we only define self-confirming solutions for the model at hand, though it seems straightforward to extend this to other robust objectives in dynamic settings with limited commitment. This definition reflects some misspecification about the commitment power of nature, thus motivating the use of the term "self-confirming." In a self-confirming solution, even though nature may have more commitment power, this extra commitment cannot hurt the seller. The seller's profit is unchanged, even if nature has more commitment power.

While self-confirming equilibria are often motivated as the solutions to a steady state learning process, in our exercise we find them intuitively appealing as the solution to the following exercise:

- A seller chooses a model of how buyers learn about their values, doing so in order to maximize their own profits.
- Upon making this choice, however, the seller reconsiders, worrying that perhaps they were wrong and not having any confidence in their understanding of the environment. The seller would be willing to abandon their model if there were some information arrival process for buyers to learn which would deliver lower profit for the seller.

A self-confirming solution—and in particular, the one we highlight—resolves the "optimism–pessimism" tradeoff highlighted by this thought experiment. An optimistic seller may assume an information structure that delivers high profits, but would reconsider this given their lack of understanding of the environment. By contrast, an overly pessimistic seller may doubt their reasons for being so pessimistic. If a price path satisfies the self-confirming criterion, a seller may think that they might as well use it, and can then rest assured that their profit guarantee would not change if in fact they were wrong.

The condition we need for the solution we highlighted to be a self-confirming one is the following:

Definition 3. *We say that a distribution F satisfies* pressed-ratio monotonicity *if $\frac{v}{F^{-1}(G(v))}$ is weakly decreasing in v.*

This assumption is satisfied for many distributions (for instance, all uniform distributions). Intuitively, the definition rules out cases where too much of the distribution is located at the top of the distribution (see also Theorem 4). In this case, a small increase threshold used in order to induce the buyer to delay leads to a larger change in the expectation of $\mathbb{E}[v \mid v \leq y]$.

Under the assumption of pressed-ratio monotonicity, we can show the following:

Theorem 3. *Suppose the value distribution satisfies pressed-ratio monotonicity. Then the equilibrium outcome in Theorem 1 is a self-confirming solution—that is, if the seller uses the outlined strategy, then there is no information arrival process which leads to lower expected payoff for the seller.*

The Theorem explicitly solves for nature's information structure under the assumption of pressed-ratio monotonicity, and shows that this involves the same information structure choice as in Theorem 1. The first step to prove this theorem is to note that the worst-case information structure is partitional. One may expect that this means the result is immediate; however, this is incorrect as our prior work (Libgober and Mu (2021)) showed via example that this property does *not* imply the worst-case information structure is the one identified in Theorem 1. That is, nature's optimal choice of information structure against a given price path *may* involve the buyer *strictly* preferring to delay purchase. Even when restricting to partitional information structures, nature's optimization problem still involves a non-trivial choice of a threshold for each time period, subject to satisfying the obedience conditions of the buyer.

We get around this issue by identifying a particular adjustment of the partition thresholds which leads to a decrease in profit whenever some threshold does not induce exact indifference when given the recommendation to not buy. While lowering the threshold induces more sale in that period, we require nature to adjust the previous period's threshold so that the buyer's indifference condition is maintained. In the Appendix, we verify that under pressed-ratio monotonicity, this will always lead to a loss of profit for the seller.

While the pressed-ratio monotonicity condition appears restrictive, we note that it will always hold in some neighborhood of the lower bound of the value distribution:

Theorem 4. *For any continuous distribution $v \sim f$ in the gap case, there exists some $y^* > \overline{v}$ such that the distribution of v conditional on being less than y^* satisfies pressed-ratio monotonicity.*

The intuition behind the theorem is simply that the pressed-ratio monotonicity condition will always hold for *generic* distributions in a neighborhood around $\underline{v}$. As a corollary, all equilibria are self-confirming solutions if the initial threshold is sufficiently close to $\underline{v}$. Alternatively, the equilibria are *eventually* self-confirming (i.e., after sufficiently many periods) if the threshold values approach $\underline{v}$, which happens whenever price discrimination becomes sufficiently fine in the limit as $\delta \to 1$.

5. CONCLUSION

This paper has grappled with the possibility of a time-inconsistent worst-case inherent in the robust approach applied to limited commitment settings. We felt our particular environment is a natural laboratory for this exercise, for various reasons; first, since the commitment solution was already solved; second, since the literature behind the Coase problem is well-explored and we were able to appeal to a wealth of intuition behind the key forces; and third, because the need to accommodate information arrival into these settings is recognized, given the past work on this and related topics.

We hope this paper has provided a template which can be used to extend the reach of the robust approach in order to obtain more insights about how limited confidence in a designer's understanding of an environment may influence their choices. On the one hand, our "as-if known values" solution in Theorem 1 seems aesthetically appealing, and appears about as simple as one could hope for as a complete equilibrium description in such settings. On the other hand, a priori it may appear at odds with our motivation of using the robust objective, since in order to obtain time-consistency we are forced to move away from allowing the seller to be concerned with all possible information arrival process. Nevertheless, this criticism, perhaps surprisingly, often turns out to have no bite. By introducing the notion of a self-confirming solution, we hope that other researchers may similarly be inspired to seek for tractable, intuitive solutions to settings with limited commitment, and can plausibly argue that they do not sacrifice anything significant behind the motivation behind their adoption of the robust approach in the first place.

References

Ausubel, L., P. Cramton, and P. Deneckere (2002): "Bargaining with incomplete information," in *Handbook of Game Theory with Economic Applications*, Elsevier, vol. 3, 1897–1945.

Ausubel, L. M. and R. Deneckere (1989): "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57, 511–531.

Bergemann, D., B. Brooks, and S. Morris (2017): "First-price auctions with general information structures: Implications for bidding and revenue," *Econometrica*, 85, 107–143.

Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

Bergemann, D. and J. Valimaki (2019): "Dynamic Mechanism Design: An Introduction," *Journal of Economic Literature*, 57, 235–274.

Bolte, L. and G. Carroll (2020): "Robust contracting under double moral hazard," *Working Paper, Stanford University and University of Toronto.*

Bose, S. and A. Daripa (2009): "Dynamic Mechanism and Surplus Extraction Under Ambiguity," *Journal of Economic Theory*, 144, 2084–2115.

Bose, S., E. Ozdenoren, and A. Pape (2006): "Optimal Auctions with Ambiguity," *Theoretical Economics*, 1, 411–438.

Bose, S. and L. Renou (2014): "Mechanism Design with Ambiguous Communication Devices," *Econometrica*, 82, 1853–1872.

Brooks, B. and S. Du (2020): "A Strong Minimax Theorem for Informationally-Robust Auction Design," *Working Paper, University of Chicago and University of California, San Diego.*

——— (2021): "Optimal Auction Design with Common Values: An Informationally Robust Approach," *Econometrica*, 89, 1313–1360.

Carroll, G. (2015): "Robustness and linear contracts," *American Economic Review*, 105, 536–563.

——— (2019): "Robustness in Mechanism Design and Contracting," *Annual Review of Economics*, 11, 139–66.

Chung, K.-S. and J. Ely (2007): "Foundations of dominant-strategy mechanisms," *Review of Economic Studies*, 74, 447–476.

Di Tillio, A., N. Kos, and M. Messner (2017): "The Design of Ambiguous Mechanisms," *Review of Economic Studies*, 84, 237–276.

Du, S. (2018): "Robust Mechanisms Under Common Valuation," *Econometrica*, 86, 1569–1588.

Duraj, J. (2020): "Bargaining with Endogenous Learning," *Working Paper, University of Pittsburgh.*

Epstein, L. and M. Schneider (2007): "Learning under ambiguity," *Review of Economic Studies*, 74, 1275–1303.

Fudenberg, D. and D. Levine (1993): "Self-Confirming Equilibrium," *Econometrica*, 61, 523–545.

Laiho, T. and J. Salmi (2020): "Coasian Dynamics and Endogenous Learning," *Working Paper, University of Oslo and University of Copenhagen.*

Libgober, J. and X. Mu (2021): "Informational Robustness in Intertemporal Pricing," *Review of Economic Studies*, 88, 1224–1252.

Lomys, N. (2018): "Learning while Bargaining: Experimentation and Coasean Dynamics," *Working Paper, Tolouse School of Economics.*

Lopomo, G., L. Rigotti, and C. Shannon (2020): "Uncertainty in Mechanism Design," *Working Paper, Duke University, University of Pittsburgh and University of California, Berkeley.*

Ocampo Diaz, S. and K. Marku (2019): "Robust contracts in common agency games," *Working Paper, Western University and University of Minnesota.*

PAVAN, A. (2017): "Dynamic Mechanism Design: Robustness and Endogenous Types," in *Advances in Economics and Econometrics*, Cambridge University Press, vol. 1, 1–62.

RAVID, D., A.-K. ROESLER, AND B. SZENTES (2020): "Learning Before Trading: On the Inefficiency of Ignoring Free Information," *Working Paper, University of Chicago, University of Toronto, and London School of Economics.*

TIROLE, J. (2016): "From Bottom of the Barrel to Cream of the Crop: Sequential Screening With Positive Selection," *Econometrica*, 84, 1291–1343.

WOLITZKY, A. (2016): "Mechanism Design with Maxmin Agents: Theory and an Application to Bilateral Trade," *Theoretical Economics*, 11, 971–1004.

A. PROOF OF THEOREM 1

To analyze this game, we first note that the buyer's problem is relatively simple. Since the buyer's decision has no effect on future prices and information (which are anyways conditional on her not purchasing), she faces an optimal stopping problem given any history.

Using the fact that we have a finite horizon, we can then turn to nature's problem and apply backwards induction. In the final period, given $(p_t)_{t=1}^{T-1}$, $(\mathcal{I}_t)_{t=1}^{T-1}$, nature chooses an information structure $\mathcal{I}_T : V \times S^{T-1} \to \Delta(S_T)$ to minimize the seller's profit. Our first goal below is to show that $\mathcal{I}_T$ can be taken to be the worst-case threshold information structure for p_2, without affecting the equilibrium outcome.

Let $s = (s_1, \dots, s_{T-1})$ be a generic signal history up until time T. Each signal history induces a posterior distribution of v, denoted F_s. First suppose s is such that the buyer does not purchase before the final period, according to the equilibrium strategy (given the price history and history of information structure, as well as the expectations of the final period prices and information). Then sequential rationality requires nature to minimize profit from this buyer type in the final period, implying that $\mathcal{I}_T(s)$ must be a worst-case information structure for the distribution F_s and price p_T. Denote the minimum value in F_s by $\underline{v}_s$, and its expected value by $\mathbb{E}[F_s]$. There are three cases:

1. If $p_T < \underline{v}_s$ or $p_T > \mathbb{E}[F_s]$, nature's problem is trivial and it is without loss to assume nature provides no information in period T.

2. If $p_T \in (\underline{v}_s,\ \mathbb{E}[F_s]]$, then for each $\epsilon > 0$, nature could reveal the worst-case threshold for $p_T - \epsilon$. This would lead to profit $p_T \cdot (1 - G_s(p_T - \epsilon))$ in period T, so equilibrium profit must be bounded above by $p_T \cdot (1 - G_s(p_T))$ by taking $\epsilon \to 0$ (note that G is continuous at p_T when $p_T > \underline{v}_s$). On the other hand, we know that equilibrium profit cannot be lower

regardless of what nature and buyer do. Hence we can without loss assume that nature provides the worst-case threshold information structure for p_T, and that *the buyer breaks indifference against the seller.*

3. The remaining possibility is $p_T = \underline{v}_s$. If F_s does not have a mass point at its lowest value, then the same argument applies since G_s is still continuous at p_T. But if F_s has a mass point of $m = G_s(p_T)$ at p_T, then any profit level in the interval $[p_T(1-m),\ p_T]$ may be supported in equilibrium, depending on how the buyer breaks ties.[9] In this case it is without loss to assume that nature reveals whether $v = p_T$ or not, and that the buyer breaks indifference in some way. Note that in this case, the seller's profit is hemicontinuous in p_T; as there is a set of possible profit levels at $p_T = \underline{v}_s$, and a unique (and continuous) profit level at $p_T < \underline{v}_s$ and $p_T > \underline{v}_s$.

Suppose instead that the signal realization s is such that the buyer purchases before the final period. In this case, we may assume nature uses the worst-case threshold information structure in the last period, which minimizes the buyer's option value (since the buyer is made indifferent between purchasing and not according to this information structure), and ensures that the buyer still purchases before the final period.

We now suppose that we have shown that nature will use a partitional information structure for all periods after the first period. We We now turn to nature's decision in period 1, showing that nature will again seek to do this in the first period. Given any price p_1 in period 1, nature expects the possibly random price $p_2 = \hat{p}_2(p_1)$ in period 2. Define the binding cutoffs w_1, w_2 by

$$w_1 - p_1 = \delta \cdot \mathbb{E}\left[(w_1 - p_2)^+\right];$$
$$w_2 = \min\{w_1, p_2\}.$$

First note that given the previous analysis, nature's information choice in period 2 leaves the buyer with the same surplus as if no information were provided in that period. Knowing this, the buyer's purchase decision in period 1 depends entirely on whether $\mathbb{E}[F_{s_1}]$ is bigger or smaller than w_1. For now, ties may be broken arbitrarily when indifferent, although we will see shortly that equilibrium requires breaking ties against the seller.

Note that, by assumption, the prior distribution F is continuous, and therefore does not have a mass point at its lowest value. We will show that nature's choice of $\mathcal{I}_1$ must be outcome-equivalent to the worst-case threshold information structure for w_1, and that the buyer must break

[9] For now we ignore the seller's optimization in the final period, and whether nature would induce such a distribution F_s in period 1. These considerations may imply that such a scenario only occurs off-path.

indifference against the seller. On the one hand, for each $\epsilon > 0$ nature could provide the threshold information structure for $w_1 - \epsilon$. Given what happens in period 2, and taking ϵ sufficiently small so that this does not influence the decision at any time after the second period, this would lead to total profit

$$p_1(1 - G(w_1 - \epsilon)) + \delta \cdot \mathbb{E}[p_2 \cdot (G(w_1 - \epsilon) - G(w_2))^+] + \sum_{s=0}^{T-2} \delta^2 p_{s+2} \mathbb{E}[p_{s+2} \cdot (G(w_{s+1}) - G(w_{s+2})^+)]$$

Letting $\epsilon \to 0$, we know that equilibrium profit following the price p_1 satisfies (taking the convention that $G(w_0) = 1$:

$$\Pi \leq \sum_{t=0}^{T} p_{t+1} \delta^t \mathbb{E}[p_t(G(w_t) - G(w_{t+1}))].$$

On the other hand, we will show that the right hand side of this expression is also a lower bound for profit, *for any choice of $\mathcal{I}_1$ and any tie-breaking rule*. Indeed, if $w_1 \leq \underline{v}$ then every type of the buyer purchases in period 1, and the result holds. Suppose $w_1 > \underline{v}$, we first show that every realization of p_2 satisfies $p_2 \leq w_1$. Recall that in period 2, any buyer who remains has expected value at most w_1. Knowing this, a price greater than w_1 leads to zero profit for the seller in period 2. This can only be optimal if the seller expects nature's equilibrium choice of $\hat{\mathcal{I}}_1$ to clear the market in period 1. We claim that this cannot occur in equilibrium. Indeed, instead of making everybody purchase, nature could reveal whether $v \in [\underline{v}, w_1)$, making this interval of buyers delay until period 2. The effect on profit is a loss of p_1 in period 1, and a gain of at most $\delta \cdot \mathbb{P}(p_2 < w_1) \cdot \mathbb{E}[p_2 \mid p_2 < w_1]$ in period 2, since these buyers purchase at p_2 only if $p_2 < w_1$. From the definition of w_1 above, we have

$$w_1 - p_1 = \delta \cdot \mathbb{P}(p_2 < w_1) \cdot \mathbb{E}[w_1 - p_2 \mid p_2 < w_1].$$

Rearranging yields $p_1 - \delta \cdot \mathbb{P}(p_2 < w_1) \cdot \mathbb{E}[p_2 \mid p_2 < w_1] = w_1 - \delta \cdot \mathbb{P}(p_2 < w_1) \cdot w_1 > 0$. Hence this deviation would lower the seller's profit.

Now that we know $p_2 \leq w_1$ almost surely, the definition of w_1 further gives $w_1 - p_1 = \delta \cdot \mathbb{E}[w_1 - p_2]$. It follows that

$$p_1 > \delta \cdot \mathbb{E}[p_2],$$

which will be useful below.

We claim that in order to minimize the seller's profit, the buyer should break ties against the seller. Indeed, the effect of delay on profit is a loss of p_1 in period 1, and a gain of at most $\delta \cdot \mathbb{E}[p_2]$

in period 2, resulting in a net decrease in profit. Next, it is without to assume nature provides only two signal realizations $\overline{s}_1$ and $\underline{s}_1$, which lead to buyer expected values $> w_1$ and $\leq w_1$, respectively. This is because any extra information in period 1 that does not change the buyer's action can be deferred to period 2. Moreover, $\underline{s}_1$ occurs with positive probability, since otherwise the market is cleared in period 1, in which case nature could deviate to lower the seller's profit as shown above.

Additionally, if $\overline{s}_1$ also occurs with positive probability, then $\underline{s}_1$ must lead to expected value exactly w_1. Otherwise, nature could mix a small fraction of $\overline{s}_1$ with $\underline{s}_1$, making this fraction of $\overline{s}_1$ no longer purchase in period 1. Suppose also that in period 2 nature separates this fraction of $\overline{s}_1$ from the $\underline{s}_1$ buyers and reveal the worst-case threshold for each group (which may not be optimal in period 2, but allows for easy comparison of profit). Then even if the fraction of $\overline{s}_1$ buyers always purchases in period 2, the profit gain is bounded above by $\delta \cdot \mathbb{E}[p_2]$. This is less than p_1, proving that the deviation would be profitable.

We can now show that the seller's profit is minimized when nature reveals the worst-case threshold for w_1 (and the buyer breaks indifference against the seller). If $w_1 \geq \mathbb{E}[v]$, then whenever $\overline{s}_1$ occurs the other signal $\underline{s}_1$ must lead to expected value less than w_1. This contradicts optimality as shown above. Thus in this case nature optimally only provides a single signal $\underline{s}_1$, corresponding to no information.

If instead $w_1 < \mathbb{E}[v]$, then $\overline{s}_1$ must occur with positive probability. So $\underline{s}_1$ leads to expected value exactly w_1. We claim that $\underline{s}_1$ must correspond to all the buyer types below the worst-case threshold for w_1. Suppose this is not the case, then we can find v' in the support of $F_{\underline{s}_1}$ and v'' in the support of $F_{\overline{s}_1}$ such that $v' > v''$. If nature were to "swap" v' and v'' with small probability, then the expected value following the modified $\overline{s}_1$ would still exceed w_1, leading to the same buyer action. Moreover, the entire posterior distribution following the modified $\underline{s}_1$ is shifted down in the FOSD sense, so profit is weakly decreased. Now since the expected value following the modified $\underline{s}_1$ is strictly less than w_1, there is room for further reducing the profit as described above. Hence the desired contradiction.

In fact, we know from this analysis that in equilibrium, nature must minimize the probability of purchase at w_1, and the buyer must break indifference against the seller. We are not done, however, since in period 1 nature could potentially provide more information than the worst-case threshold (for example making the buyer's posterior distribution supported on only two values). This would affect the seller's belief about the buyer's value distribution in period 2, and influence the optimal price p_2.

To address this issue, we are going to show that the price p_2 would remain optimal if nature were to simply provide the worst-case threshold information structure for w_1 in period 1. To this end, note that in this equilibrium, any realization of p_2 must be maxmin optimal against a buyer who knows her value to be in the lowest $G(w_1)$-percentile and potentially knows more.

Moreover, as calculated above, the maxmin optimal profit in period 2 must be $p_2(G(w_1) - G(w_2))$ (which must be the same number for all realizations of p_2). Now, against a less informed buyer who only knows her value to be below the $G(w_1)$-percentile, the maximal optimal profit can only decrease. But charging price p_2 against such a buyer guarantees $p_2(G(w_1) - G(w_2))$, so it remains the seller's best response.

Hence, we have shown that every equilibrium is outcome-equivalent to an equilibrium in which nature provides threshold information structures, where the threshold is chosen so that conditional on having value below the threshold, the buyer is indifferent between purchasing in the current period or delaying until the future (without further information). Moreover, the seller thinks the buyer always breaks indifference against him (even though this is not necessarily true in period 2, if nature has deviated in period 1). Therefore, given any equilibrium price path shaping expectations, the seller's probability of sale in each period under any deviation strategy is the same as the known-values case, with G replacing F as the value distribution. It follows that any equilibrium in our model is equivalent to an equilibrium in the known-values case with the transformed value distribution G.

B. PROOFS FOR SECTION 3

Proof of Proposition 1. Fix a distribution of value F and any time horizon K. Let $\mathcal{U}(F)$ be the expected value of the buyer under known values when their value is drawn according to F. The idea behind the construction is the following:

- The buyer learns their true value at time K, and the equilibrium played is the unique equilibrium for the gap case.
- Before time K, the seller chooses a price path which makes the buyer indifferent between purchasing before time K and waiting until time K (at which point they will obtain $\mathcal{U}(F)$).
- Meanwhile, the buyer randomizes purchase so that the seller has incentives to follow the equilibrium strategy.
- If the seller deviates, the equilibrium reverts to the worst-case outcome outlined in Theorem 1.

We now walk through the details more precisely. The price path the seller charges is:

- $\mathbb{E}[v] - \delta^{K-t}\mathcal{U}(F)$ for all $t < K$

- $\hat{p}_s(F)$ in period $K+s-1$, where $\hat{p}_s(F)$ is the equilibrium price charged in the unique gap equilibrium under known values with distribution F.

The information structure available to the buyer is simple: On path, no information is provided until period K, but in period K, the buyer learns the value perfectly.

Any deviation (by any part) at or after time K is treated in exactly the same way as in the known-values case. If the seller deviates to some price p' *before* time K, then the buyer obtains a partitional information structure, where the threshold is set to be $F^{-1}(G(\hat{w}(p')))$, where $\hat{w}(p')$ is the value at which the buyer would be indifferent between purchasing and not assuming the seller followed the equilibrium strategy from that time on.

It remains to specify buyer behavior before time K. Note that, since $\delta < 1$ and the equilibrium strategy involves delay, we have $\mathbb{E}[v] > \pi(F) + \mathcal{U}(F) > \pi(F) + \delta\mathcal{U}(F)$. But we also have $\pi(F) > \pi(G)$. Therefore, we specify that the buyer purchases in period $K-1$ with probability $\rho_{K-1} < 1$, where ρ_{K-1} satisfies:

$$\pi(G) < \rho_{K-1}(\mathbb{E}[v] - \delta\mathcal{U}(F)) + (1-\rho_{K-1})\delta\pi(F).$$

Indeed, while $\delta\pi(F) < \pi(G)$ may hold, choosing ρ_{K-1} sufficiently close to 1 means the seller obtains a larger payoff by choosing the price in that period, $\mathbb{E}[v] - \delta\mathcal{U}(F)$, than deviating.

Now suppose we have chosen $\rho_{K-1}, \ldots, \rho_{K-n} < 1$ so that the seller obtains a continuation value larger than $\pi(G)$ by choosing the specified prices at times $K-n, \ldots, K-1$ (and after, once the buyer knows her value); call this value V_{K-n}. We again have $\mathbb{E}[v] - \delta^{n+1}\mathcal{U}(F) > \pi(F) > \pi(G)$, We can then specify that the buyer purchases in period $K-n-1$ with probability $\rho_{K-n-1} < 1$, where ρ_{K-n-1} is such that:

$$\pi(G) < \rho_{K-n-1}(\mathbb{E}[v] - \delta^{n+1}\mathcal{U}(F)) + (1-\rho_{K-n-1})\delta V_{N-k}.$$

Again, since the buyer's expected payoff from time K (once they know their value is $\mathcal{U}(F)$, they are indifferent between delaying until time K and following their optimal strategy, and buying at any time prior to time K. Therefore, they are willing to follow this mixed strategy. If the seller deviates, however, then immediately their payoff is bounded above by $\pi(G)$. By construction, the purchase probabilities when the seller follows the equilibrium strategy are chosen such that the seller's payoff is higher by following the equilibrium. We therefore have that at every time, the seller finds it optimal to choose the specified price, and the buyer finds it optimal to follow the specified strategy, proving the proposition. □

C. PROOFS FOR SECTION 4

Details for Example 1. We perform the familiar calculation for the equilibrium price path by backwards induction using this known values distribution, using the fact that the equilibrium is of a threshold form. First, note that given an arbitrary first period indifference threshold $\overline{v}$ under known values, we have the seller's second period price must maximize $p_2(1-\frac{p_2}{\overline{v}})$, implying that $p_2=\frac{\overline{v}}{2}$. Anticipating this and observing a first period price of p_1, the buyer is indifferent if:

$$\overline{v}-p_1=\delta\left(\overline{v}-\frac{\overline{v}}{2}\right)\Rightarrow\overline{v}=\frac{2p_1}{2-\delta}.$$

Therefore, the seller at time 1 choose p_1 to maximize:

$$p_1\left(1-\frac{2p_1}{2-\delta}\right)+\delta\frac{p_1}{2-\delta}\left(\frac{p_1}{2-\delta}\right)\Rightarrow 1-\frac{4p_1}{2-\delta}+\frac{\delta 2p_1}{(2-\delta)^2}=0\Rightarrow p_1=\frac{(2-\delta)^2}{8-6\delta}.$$

Substituting this in gives that profit is:

$$\begin{aligned}\frac{(2-\delta)^2}{8-6\delta}\left(1-\frac{2-\delta}{4-3\delta}\right)+\delta\frac{(2-\delta)^2}{(8-6\delta)^2}&=\frac{(2-\delta)^2}{8-6\delta}\left(1-\frac{2-\delta}{4-3\delta}+\frac{\delta}{8-6\delta}\right)\\&=\frac{(2-\delta)^2}{8-6\delta}\left(\frac{4-3\delta}{8-6\delta}\right)=\boxed{\frac{(2-\delta)^2}{4(4-3\delta)}}.\end{aligned}$$

Now we compute the profit under the information structure specified in Theorem 2. First, recall that $\pi^*(\tilde{v})=\frac{\tilde{v}}{8}$. Since $\mathbb{E}[v\mid v<\tilde{v}]=\tilde{v}/2$, the buyer obtains $\frac{3\tilde{v}}{8}$ in the second period. Therefore, the buyer's continuation value, given $\tilde{v}$, solves:

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

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. Let us assume for the moment that this solution involves purchase in each period with positive probability, handling the case where this does not occur separately. 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:

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

Profit at this price is:

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

Unlike with the previous case, however, we need to check that this solution does indeed involve sale at both periods. 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 $\delta/4$.

Proof of Theorem 2. Let $p_1>p_2>\ldots>p_{t^*}=\underline{v}$ be a solution to the baseline model, with corresponding thresholds $y_1>y_2>\cdots>y_{t^*}=\underline{v}$. Let U_2 denote the buyer's expected continuation surplus in this equilibrium starting at the second period, and let Π_2 denote the seller's continuation profit. Note that:

$$\int_{\underline{v}}^{y_2}wf(w)dw>U_2+\Pi_2,$$

since by assumption the baseline model does not involve the market clearing by time two. The idea is to use the fact that there is inefficiency to transfer additional surplus to the buyer in order to induce additional delay.

We do this by considering the following classes of information structures for nature:

- In period 1, nature chooses a threshold $\tilde{y}_1$ as a function of the first period price, the seller charges.
- In the second period, if the seller chooses some fixed $p_2=\tilde{\Pi}$, then nature reveals no information to the buyer, and reveals no information to the buyer in the future.
- If the seller uses some other price, nature uses the worst-case descending partitional information structure outlined in the proof of Theorem 1.

We will in particular focus on the case where $\tilde{\Pi}$ is the seller's continuation profit follow some first period threshold of y_1, which we denote $\Pi_2(y_1)$. Note that in this case, the seller has a best reply to choose $p_2 = \Pi_2(y_1)$, since by construction deviating cannot lead to a higher profit (otherwise, there would be some other strategy yielding higher profit in the baseline model).

Now, nature choosing some information structure of this form may induce the seller to choose a price such that the market would clear at time 1 or time 2. However, the seller also had the ability to charge one of these prices in the baseline model, and did not, meaning that this will hurt the seller.

On the other hand, for any other price, we have that the threshold y_1 such that the buyer is willing to not purchase whenever informed that their value is below the threshold satisfies $y_1 > F^{-1}(G(p_1))$, since, by the previous, their continuation surplus increases. It follows that under this class of information structures, the seller sells less in the first period relative to the case without nature commitment, and obtains the same continuation profit, and therefore obtains lower discounted expected profit, as desired. □

Proof of Theorem 3. We fix an arbitrary declining price path $p_1, \ldots, p_{t^*}$ with $p_{t^*} = \underline{v}$. We note that in the gap case, such a t^* exists for every equilibrium price path whenever $\delta < 1$ under a known value distribution. Therefore, using the previous result, such a t^* can be always be found in any equilibrium of the game without nature commitment. Furthermore, by Proposition 3 in Libgober and Mu (2021), the worst-case information structure against an arbitrary declining price path is a threshold process. It follows that nature's choice of information structure is determined by thresholds $y_1 > y_2 > \cdots > y_{t^*} = \underline{v}$, with the buyer purchasing at the first time t satisfying $v > y_t$.

We first note that the buyer always purchases at or before period t^*. The theorem will follow from showing that each threshold y_t should be as low as possible, for all $t < t^*$. For the first part of the proof, we consider any information structure with $y_1 > y_2 > \cdots > y_{t^*}$; we address the case where equality might hold separately. That is, we show that a buyer who does not purchase at some time t must be *indifferent* between purchasing and continuing in any worst case information structure. This is immediate for y_1; In this case, increasing y_1 while holding all other thresholds fixed simply trades off between sale at time 1 and time 2; so, if y_1 could be raised without changing the buyer's incentive conditions, since $p_1 > \delta p_2$, this hurts the seller.

Suppose we have that y_t is set so that the buyer is indifferent between purchasing and continuing when given the recommendation to not purchase. This gives us the following indifference condition, given our threshold sequence:

$$\int_{\underline{v}}^{y_t} (v - p_t) f(v) dv = \sum_{s=t+1}^{t^*} \delta^{s-t} \left(\int_{y_s}^{y_{s-1}} (v - p_s) f(v) dv \right). \tag{2}$$

In addition, we have the following expression for the seller's profit, using the convention that $F(y_0) = \overline{v}$:

$$\sum_{s=1}^{t^*} p_s (F(y_{s-1}) - F(y_s)). \tag{3}$$

We will prove that, under the assumption of pressed-ratio monotonicity, if y_{t+1} does not induce the buyer to be indifferent between purchasing and continuing at time $t+1$ (i.e., if the buyer strictly prefers to continue), then the thresholds can be adjusted to lower the seller's profit.[10] In particular, we will show that if nature adjusts y_t to maintain the buyer's indifference at time t between purchasing and continuing, then lowering y_t will increase profit.

Under this particular perturbation, we can differentiate (3) with respect to y_{t+1}, using (2) to implicitly differentiate $y_t(y_{t+1})$. The derivative of the right hand side of (3) with respect to y_{t+1}, holding fixed y_s for $s > t+1$, is:

$$\delta(-(y_{t+1} - p_{t+1}) + \delta(y_{t+1} - p_{t+2}))f(y_2).$$

Let $(1-\delta)\overline{v}_{t+1} = p_{t+1} - \delta p_{t+1}$, so that $\overline{v}_{t+1}$ is indifferent between purchasing and continuing at time $t+1$, and rewrite the derivative of the right hand side as:

$$\delta(1-\delta)(\overline{v}_{t+1} - y_{t+1})f(y_2).$$

We note that this derivative is negative as long as $y_{t+1} > \overline{v}_{t+1}$. Hence decreasing y_{t+1} increases the value of the right hand side, whenever y_{t+1} is above the indifferent value. We now differentiate the indifference condition with respect to y_t, after the term on the right hand side of (2) involving y_t is added to the left hand side:

$$(y_t - p_t)f(y_t) - \delta(y_t - p_{t+1})f(y_t) = (1-\delta)(y_t - \overline{v}_t)f(y_t),$$

with $\overline{v}_t$ defined analogously. Thus, our previous work together with chain rule implies:

$$\delta(\overline{v}_{t+1} - y_t)f(y_{t+1}) = (y_t - \overline{v}_t)f(y_t)y_t'(y_{t+1}). \tag{4}$$

[10]To emphasize, by itself, decreasing y_{t+1} will increase the seller's profit, by inducing more sale at time $t+1$, as opposed to late times where the seller obtains less.

Note that since $y_{t+1} > \overline{v}_{t+1}$ and $y_t > \overline{v}_t$, we have $y_t'(y_{t+1}) < 0$; thus lowering the time $t+1$ threshold decreases the probability of sale at time t. The observation that $y_t'(y_{t+1}) < 0$ will be useful later in the proof.

We are now ready to differentiate (2). Under the particular perturbation listed, since only y_t and y_{t+1} adjust, we have it suffices to differentiate:

$$p_t(1 - F(y_t(y_{t+1}))) + \delta p_{t+1}(F(y_t(y_{t+1})) - F(y_{t+1})) + \delta^2 p_{t+2}F(y_{t+1}),$$

as all other terms are constant. Differentiating yields:

$$-p_t f(y_t(y_{t+1}))y_t'(y_{t+1}) + \delta p_{t+1}(f(y_t(y_{t+1}))y_t'(y_{t+1}) - f(y_{t+1})) + \delta^2 p_{t+2} f(y_{t+1}).$$

Now, multiply through by $(y_t - \overline{v}_t)$ (which we recall is positive), and use (4) to eliminate the right hand side wherever it appears in the derivative of profit with respect to y_{t+1}; doing this and factoring out terms, we have that the derivative of profit with respect to y_{t+1} is proportional to:

$$\delta f(y_{t+1}) \cdot (-(p_t - \delta p_{t+1})(\overline{v}_{t+1} - y_{t+1}) - (p_{t+1} - \delta p_{t+2})(y_t - \overline{v}_t)).$$

To find the change in profit from *lowering* y_{t+1} (as opposed to raising it), we must multiply this by -1. Doing this, and substituting in for $\overline{v}_t$ and $\overline{v}_{t+1}$, we have the change in profit from lowering the y_{t+1} threshold (and hence departing from the "known but pressed" outcome) is proportional to:

$$\overline{v}_t(\overline{v}_{t+1} - y_{t+1}) + \overline{v}_{t+1}(y_t - \overline{v}_t) = -\overline{v}_t y_{t+1} + \overline{v}_{t+1} y_t. \tag{5}$$

Note that, by the pressed-ratio monotonicty assumption, this expression is positive when y_{t+1} satisfies $\mathbb{E}[v \mid v \leq y_{t+1}] = \overline{v}_{t+1}$ (i.e., the value corresponding to the pressed threshold), which is exactly when y_{t+1} is as large as possible. It follows that, when y_t is chosen so that this equation holds with equality, profit is locally increasing if y_t is lowered.

On the other hand, suppose y_{t+1} is lower than the threshold inducing the pressed distribution. Note that nowhere in the above derivation, except when we signed the derivative, did we use that y_{t+1} was set to be the threshold corresponding to the pressed distribution. Now, notice that if we multiply the right hand side of (5) by -1 and differentiate, we have:

$$\overline{v}_t - \overline{v}_{t+1} y_t'(y_{t+1}) > 0.$$

This implies that the right hand side of (5) is actually *smallest* when y_{t+1} is as large as possible.

Since it is positive at this value, this means that it is positive everywhere. While this does not imply profit is convex in y_t (since profit depends on $\delta f(y_{t+1})$, which we have dropped), it does imply that (5) is positive for *all* choices of y_{t+1} in the relevant range. In other words, this shows that nature can always decrease profit by increasing y_{t+1} according to this perturbation.

We have therefore shown that any partitional information structure with thresholds $y_1 > y_2 > \cdots > y_{t^*}$ can be made worse for the seller if there is some period where the buyer strictly prefers to delay purchase, given the anticipated price path. It remains to consider the case where some thresholds may hold with equality. Suppose $y_s = y_{s+1} = \cdots = y_{s+k}$. There are two cases to consider:

- Lowering all thresholds simultaneously does not lead to a violation of the obedience constraint. In this case, the argument is identical, simply by collapsing all periods where trade does not occur into a single period.
- Lowering all thresholds simultaneously leads to the obedience constraint being violated for period s. In that case, the same argument implies keeping the thresholds at time $s+1, \ldots, s+k$ holding with equality while rising the threshold at time s would lower the seller's profit.

That these are the only two cases to consider follows from the fact that the thresholds are declining over time. This proves the theorem. □

Proof of Theorem 4. We consider the derivative of $\frac{v}{F^{-1}(G(v))}$:

$$\frac{d}{dv}\frac{v}{F^{-1}(G(v))} \propto F^{-1}(G(v)) - v\frac{d}{dv}F^{-1}(G(v)).$$

Also recall that $F^{-1}(G(v)) = L^{-1}(v)$, where $L(y) = \mathbb{E}[v \mid v \leq y]$. By the inverse function theorem, we differentiate L^{-1} as follows:

$$\left.\frac{d}{dv}F^{-1}(G(v))\right|_{v=\tilde{v}} = \frac{1}{L'(y)},$$

where y is the threshold that leads to $\mathbb{E}[v \mid v \leq y] = \tilde{v}$. As will become important later, we note that $\lim_{\tilde{v}\to\underline{v}} L^{-1}(\tilde{v}) = \underline{v}$.

Since $L(y) = \frac{\int_{\underline{v}}^{y} wf(w)dw}{F(y)}$, we can differentiate the function $L(y)$ as follows:

$$L'(y) = \frac{f(y)\left(yF(y) - \left(\int_{\underline{v}}^{y} wf(w)dw\right)\right)}{F(y)^2}.$$

We note that this function shares the same differentiability properties as F whenever $y > \underline{v}$. In order to prove the theorem, we study the limit of this expression as $y \to \underline{v}$. Notice that in the limit as $y \to \underline{v}$, both the numerator and the denominator approach 0. By L'Hopital's rule, however, to evaluate this limit, we can differentiate the numerator and the denominator twice to obtain:

$$\lim_{y \to \underline{v}} L^{-1}(y) = \lim_{y \to \underline{v}} \quad \frac{(f(y))^2 + 2F(y)f'(y) + (yF(y) - \int_{\underline{v}}^{y} wf(w)dw)f''(y)}{2(f(y)^2 + F(y)f'(y))}.$$

However, since $F(\underline{v}) = 0$, we have that this limit reduces very simply to $\frac{1}{2}$.

Returning to the original limit, and recalling that $\lim_{\tilde{v} \to v} F^{-1}(G(v)) = \tilde{v}$, we therefore put this together to obtain the following:

$$\lim_{\tilde{v} \to \underline{v}} \quad \frac{d}{dv} \frac{v}{F^{-1}(G(v))}\bigg|_{v=\tilde{v}} = \underline{v} - \underline{v}\frac{1}{1/2} = -\underline{v} < 0.$$

Using the differentiability properties of the distribution, we therefore have that pressed-ratio monotonicity condition is satisfied in some neighborhood of $\underline{v}$, as desired. $\square$