Coasian Dynamics under Informational Robustness

Jonathan Libgober* and Xiaosheng Mu**

*University of Southern California
**Princeton University

April 15, 2022

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 introduce a formulation whereby the seller chooses prices to maximize profit against a *dynamically-consistent* worst-case information structure. In the gap case, 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, allowing for arbitrary (possibly dynamically-inconsistent) worst-case 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 *reinforcing solution.* As other formulations of our problem introduce dynamic-inconsistency, the notion of a reinforcing solution may be useful for researchers seeking to tractably relax the commitment assumption while maintaining a robust objective. To highlight the non-triviality of these observations, we show that while the analogy to known values can hold under an equilibrium selection in the no-gap case, it does not hold more generally.

Keywords. Durable goods monopoly, limited commitment, dynamic informational robustness, reinforcing equilibrium.

Contact. libgober@usc.edu, xmu@princeton.edu. We thank Nageeb Ali, Gabriel Carroll, Andrew Ellis, Mira Frick, Drew Fudenberg, Kevin He, Ryota Iijima, Yuhta Ishii, Navin Kartik, Andreas Kleiner, Francesco Nava, João Ramos and Balazs Szentes for helpful comments.

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, were the buyer perfectly informed of their willingness-to-pay, the monopolist should be expected to obtain arbitrarily low levels of profit. Intuitively, a monopolist charging a high price would be expected to cannibalize their future demand; as a result, a buyer anticipates future price drops, creating an additional incentive to delay purchase. The monopolist would respond to the possibility of delay by lowering the price in the initial period; provided players are sufficiently patient, this process unravels to the point where the monopolist is left charging a price close to the lowest possible buyer value.

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 of exogeneity of buyer information over willingness-to-pay is one of them. If the buyer does not exactly know their value for the seller's product, then the possibility emerges that information might only be decision-relevant were the seller to choose certain prices. In this paper, we note that, if the buyer's scheme for acquiring information were necessarily known to the seller, then this feature could completely reverse the Coase conjecture (Proposition 2). 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 learning process, it is not obvious a priori which constraints on information arrival can be imposed without ruling out economically meaningful possibilities.[1]

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 such concerns can justify mechanisms much simpler than those seeking to optimize against the precise details of

[1]One such restriction would be to assume information is independent of prices. While a sensible assumption when the seller only moves ones, this appears much less plausible when the seller re-optimizes in every period. It seems natural to posit that, for instance, that buyers might pay more attention to the product if the seller does something unexpected; this kind of reaction would be ruled out if we assumed information were independent of prices.

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 its applicability.

The key 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 still typically require some added structure to analyze fully,[2] (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 not (yet) 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 moving away from strong 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

[2] 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 in 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 past work, specifically Libgober and Mu (2021).

In this paper, we primarily focus on the gap case,[3] where the Coase conjecture obtains in the unique equilibrium in the absence of learning (but, as mentioned above, not necessarily in its presence); we comment on the no-gap case below and more patiently in Section 5. Our first contribution is to identify a way of specifying the robust objective under limited commitment whereby the worst-case is *dynamically-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[4] 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, the commitment assumption plays no role. 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 defined as *pressing*). By contrast, this paper shows that, provided the information structure at every period is chosen to minimize the continuation payoff from that time on, the key properties of the solution are maintained, despite some added technical details.

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

[3]That is, where there is a gap between the seller's cost and the lowest possible buyer value

[4]Part of the reason we focus primarily on the "gap" case is to avoid conceptual difficulties which emerge when the worst-case depends on equilibrium selection. In the no-gap case, the worst-case profit will not be uniquely pinned down, since multiple payoff levels can be achieved depending on which equilibrium strategy the buyer follows.

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 *reinforcing solution*, and show that it is often satisfied by the benchmark solution outlined above. In a reinforcing 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* against *any* arbitrary dynamic information arrival process.

The sufficient condition we identify for our baseline solution to be a reinforcing 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 more patient than sellers), then our baseline model's solution will *always* be reinforcing.

We briefly comment that the characterization might not be as sharp when there is no gap. Under no gap and known values, the Coase conjecture need not emerge in every equilibrium. In this paper, we further show that the added richness in equilibrium outcomes under known values can be used to sustain other equilibria where the outcome is not analogous to *any* known values environment. This highlights the non-triviality of our results, and raises many natural questions for follow-on work regarding which other assumptions on a designer's problem could single out unique (and sensible) worst-case solutions.

Nevertheless, since we do obtain unique, intuitive solutions under some economically reasonable assumptions, our main message in many cases, the dynamic inconsistency for robust limited commitment objectives issue may be less severe than originally thought—though this should not be taken for granted. 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 believe this solution concept could be used elsewhere. Notice that a correctly specified Bayesian decisionmaker would vacuously choose a reinforcing solution. So, a natural question is whether other interesting mechanism design settings also possess intuitive reinforcing 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 by the goal of relaxing 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), subsequent literature considered the case of "unknown values" case where the designer also faced uncertainty about what participants knew about their own preferences (Bergemann et al. (2017), Du (2018), Brooks and Du (2021), Brooks and Du (2020), Libgober and Mu (2021)).[5] 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 at first blush. 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.[6] 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.[7] 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 view as convincing). As far as we are aware, the main difficulty in relaxing commitment 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

[5]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.

[6]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.

[7]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.

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 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 does not arise directly.

A less related literature considers mechanism design 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 dynamic-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 under commitment (see Appendix F of Libgober and Mu (2021)), the "Bayesian agent" case is worst-case for the seller (at least under some forms of ambiguity aversion).

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 departures might emerge is that learning can change the direction and magnitude of selection pressures, both of which are crucial for the Coasian dynamics to emerge (see Tirole (2016)). We show that such forces are essentially absent in the robust case—at least, in the gap case, and in at least one equilibrium in the no-gap case.

2. MODEL

A seller of a durable good interacts with a buyer in discrete time until some terminal date T, where $T \leq \infty$, though we will handle the case of $T = \infty$ and $T < \infty$ separately. 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. We comment on the no-gap case only in Section 5.

Within a period t, the seller chooses a price $p_t \in \mathbb{R}_+$, after which the buyer decides whether to purchase or not. 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. In principle, the seller can choose a randomization over p_t, say $\gamma_t \in \Delta(\mathbb{R}_+)$ but the realization of this randomization will be known to nature prior to its choice.

- 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). We assume this information structure is observed both by the seller and the buyer.

- 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 $-\delta^{t-1}p_t$). 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:

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

By contrast, at each time t, nature chooses the information structure to maximize:

$$-p_t\mathbb{P}[\sigma_{s^{t-1},I^t}(s_t) = 1 \mid p_1, \ldots, p_t, \sigma] - \sum_{k=t+1}^{T} \delta^{k-t+1}\mathbb{E}_{p_k \sim \gamma^k}[p_k\mathbb{P}[\sigma_{s^{k-1},I^k}(s_k) = 1 \mid p_1, \ldots, p_k], \quad (2)$$

i.e., the negative of (1). Note that, even though nature maximizes the negative of the seller's payoff (and visa versa), since this is a three player interaction (due to the presence of the buyer), strictly speaking this is not a zero-sum game.

2.3. Discussion

Our explicit use of "nature" as a player is primarily expositional device to explain why one might expect dynamic- 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.

Our assumption that the seller observes nature's choice reflects a strong pessimism that the information structure is *actually* chosen in order to minimize the seller's profit. Note that when this choice is observed, nature always has the option of choosing the information structure they would have picked were it not observed. Therefore, the choice of making the information structure observable reflects one of the subtleties of specifying the robust objective under limited commitment. This reflects a high degree of confidence in pessimism from the seller. An interesting question, which we address in part by studying reinforcing solutions below, is whether the assumption of pessimism would change the outcomes that one might expect to emerge.

One issue that emerges more generally in dynamic models under a robust objective is how the timing of nature's moves interactions with the individual seeking robustness. While we allow the seller to randomize the price in every period, we also allow the information structure in a given period to depend on that price. Our view is that this is the most natural benchmark given the focus on limited commitment. The reason is that with limited commitment, the seller seeks robustness *at every time*. Therefore, it is natural to think that the information the seller is concerned about at time 10 may depend on their actions at time 5, immediately imposing at least some degree of dependence of information on the price. Furthermore, to assume that the information structure did not react to the seller's choice in that period would imply, for instance, that information could not react to a deviation for the seller. But it is easy to imagine that consumers might receive some different information if a seller acts unexpectedly. Therefore, while alternatives may be interesting, our view is that this timing protocol is particularly natural for a durable goods seller under limited commitment (more so than in the case of commitment, where—in contrast to this paper—the seller's actions are necessarily before *any* moves of nature).

A natural question relates to 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 Libgober and Mu (2021), which showed the optimal selling strategy is a constant price path.[8] More precisely, that paper identified a *known-values* environment which would deliver an identical pricing strategy as optimal. Despite the similarity,

[8]If the seller commits to a pricing strategy, then whether nature has commitment does not change the solution whenever a constant price path emerges. The reason is that in those cases, it turns out the worst-case involves all information being revealed in the initial period. So, if the buyer does not buy immediately, then she never will and hence nature could not gain by changing their actions in future periods.

the conceptual difficulties are 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 $T = 1$.

3. SOLUTION TO THE BASELINE MODEL

We now proceed to solve the previous model.[9] 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, 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, Libgober and Mu (2021) also showed by example that one should generally not expect the pressed distribution to characterize the seller's problem if a declining price path were used. 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 here given that we are focused on the noncommitment case (where prices decline).

Our first result shows that those information structures are dynamically-inconsistent, in that they rely upon giving the buyer more information than the worst-case at later times. If one forces

[9] 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. See Section 5 for more on the no-gap case.

those information structures to also minimize the seller's profit from that time on, then we again recover the tight analogy:

Theorem 1. *When $T < \infty$, 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 as follows. 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.

Due to our focus on the gap case, we can also show the following:

Proposition 1. *Suppose the distribution F involves $\underline{v} > 0$ and satisfies the Lipschitz condition of Theorem 4 of Ausubel et al. (2002).*[10] *When $T = \infty$, there exists some finite period $\hat{T}$ such that the market clears by $\hat{T}$; therefore, the same conclusion from Theorem 1 holds when $T = \infty$.*

This result uses the fact that the equilibrium outcome under known values features a finite horizon. In our problem, *if* the outcome were that of Theorem 1, then we would have the same

[10] In our notation, this requires that

$$F^{-1}(q) - \underline{v} \leq Lq,$$

for some L and all $q \in [0, 1]$.

objective defining the seller's objective. The difficult part is showing that this is in fact all that can happen. That the seller has no profitable deviation, if information is chosen to minimize their total discounted payoff at every period, is fairly straightforward, since this is true under known-values, and hence true even if nature only uses partitional information arrival processes. The argument for nature is that, for *any* candidate equilibrium information structure, the best-case reaction from buyers *for the seller* would be to assume no further information were received. Therefore, to derive an upper bound on the seller's equilibrium profit (i.e., ask "how badly can nature possibly do?"), it is enough to assume that this is the inference the buyer would make following a deviation of nature. Thus, the highest profit the seller could obtain in a given period does not necessarily depend on future information structure choices, allowing us to derive an upper bound on the equilibrium profit. Noting that this coincides with the value function assuming the partitional equilibrium, we then conclude the worst-case information structure is again essentially unique (i.e., induces a unique response from the buyer).

Theorem 1 and Proposition 1 provide 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. This sharpness should not be taken for granted. The proof of Proposition 1 uses the result that under known values, there are a finite number of periods after which the market clears (stated in Ausubel et al. (2002)). 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 2. *Fix F, δ and T. Suppose the equilibrium outcome under known values with distribution 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 seller uses a constant price path.*
- *The seller obtains continuation value of v^* at every point in time, where v^* is less than $\mathbb{E}_F[v]$ but larger than the minimax profit from Theorem 1 given any time horizon $k \leq T$.*
- *The market does not clear in any finite time.*

The clearest failure of Coasian forces identified in the previous Proposition is the impossibility of bound the time at which the market will clear over the set of all information structures. The other two could arguably be obtained under specific known-values environments, where the market clears at time 1. 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. Even in the no-gap case, we are not aware of any model which delivers the market failing to clear in finite time and the seller using a constant price path, which the equilibrium constructed in 2 features. 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.

Looking ahead, it turns out that when there is no gap, such equilibria may emerge even in our baseline model (though restricting buyer behavior to minimize seller profit would rule them out). As a result, the analogy to known-values *requires* the no-gap assumption. This contrasts with the case where the seller has commitment, where no such qualifiers emerge.

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 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 there exists an information arrival process and sequential equilibrium such that the seller obtains a lower expected profit than in the unique equilibrium outlined in Theorem 1.*

The following example illustrates:

Example 1. *Suppose $T = 2$ and $v \sim U[0, 2]$. 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 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*

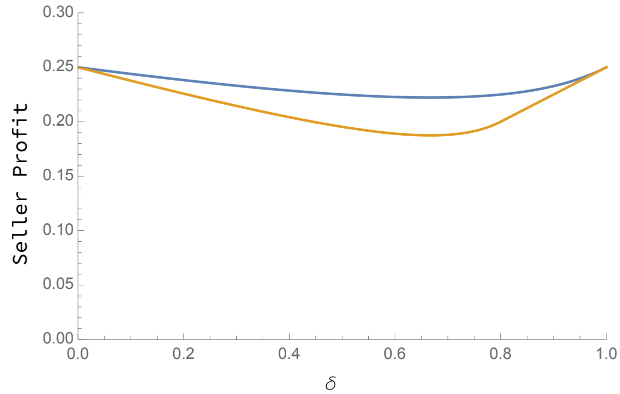

Figure 1: Comparison of seller profit between equilibrium of the baseline model vs. Theorem 2, for Example 1

$\delta \geq 4/5$ (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 (orange line). 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 **reinforcing 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 any similar concept being studied elsewhere in the robust mechanism design literature, though we view it as very natural. To maintain focus, we only define reinforcing 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. In a reinforcing 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.

In our exercise, we find reinforcing solutions intuitively appealing as the solution to the following exercise:

- A seller chooses a model of how buyers learn about their values, doing so in an optimistic way in order to *maximize* their own profits.
- Upon making this choice, however, the seller becomes pessimistic and reconsiders; the worry is that perhaps they were wrong, and they also lack confidence in their understanding of the environment. Therefore, the seller would abandon a model if there were *some* information arrival process the buyer could have which would deliver lower expected profit for the seller.

A reinforcing 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 reinforcing 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—no matter how pessimistic they are.

The condition we need for the solution we highlighted to be a reinforcing 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 Proposition 3). 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 reinforcing 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 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:

Proposition 3. *For any continuous distribution $v \sim f$ in the gap case, there exists some $y^* > \underline{v}$ such that the distribution of v conditional on being less than y^* satisfies pressed-ratio monotonicity.*

As a corollary of this proposition, all equilibria are reinforcing solutions if the initial threshold is sufficiently close to $\underline{v}$. Alternatively, the equilibria are *eventually* reinforcing (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. THE NO-GAP CASE

Our analysis so far has assumed that $\underline{v} > 0$, which past work has shown is a key assumption to deliver the Coase conjecture under known values. We note that, in the case of a finite horizon, identical results apply to the no-gap case as well. However, with an infinite horizon, the story is different. On the one hand, Ausubel and Deneckere (1989) show that in the no-gap case, an equilibrium exists ensuring that the monopolist obtains arbitrarily low levels of profit as the time between offers shrinks to 0. Though trade does not end in finite time (since values being arbitrarily close to 0 means that, as long as the monopolist does not charge 0, the market is not cleared), this equilibrium is otherwise Coasian, as the market anticipates that the monopolist will cannibalize future demand. Using this equilibrium, however, they are able to derive a folk theorem which ensures that the monopolist obtains a profit level very close to what would be obtained under commitment. The idea is simple: A monopolist is deterred from lowering prices too much, at every point in time, via a punishment which reverts to the Coasian equilibrium where profit levels are arbitrarily low.

The lack of a gap does not change the construction of a worst-case equilibrium given a declining price path. Therefore, we can obtain the same analogy to the known-values case, and obtain a similar kind of folk theorem as in Ausubel and Deneckere (1989). However, in light of Proposition 2, one may wonder whether new kinds of equilibria, fundamentally different from those under known values, might arise in the no-gap case. As we have emphasized, when there *is* a gap, there is essentially no difference.

However, since the folk theorem of Ausubel and Deneckere (1989) shows a *range* of possible outcomes for the seller, we can use their constructions to not only discipline the behavior of the seller, but nature as well. In this context, not only does this enable the possibility of an indeterminacy in the seller's payoff, but also that the corresponding outcome may be qualitatively different from any equilibrium under known values, thus also dramatically breaking the analogy between the two settings.

Proposition 4. *Suppose $\underline{v} = 0$, and that the distributions F and G satisfy Definition 5.1 of Ausubel and Deneckere (1989). Then the information structure from Proposition 2 can emerge as an equilibrium outcome.*

The proposition is noteworthy because not only does it demonstrate that in the gap case we may have a failure of the Coase conjecture, *but also a failure of the analogy to known-values.* The equilibrium described in Proposition 2 is unlike any that can emerge under known (and constant) values, since (a) the buyer obtains zero surplus and yet (b) the market never clears. It is worth noting that subtleties such as these fail to emerge in the commitment case. There, the uniqueness is much more immediate, since the seller essentially faces a decision problem, only taking an action once before anyone else. However, the fact that the limited commitment setting is necessarily a game means such uniqueness can no longer be taken for granted; and indeed, once uniqueness fails, so too might the analogy to known values.

Of course, if the equilibrium selection were chosen to minimize the seller's profit, then these issues would not arise and the equilibrium would still feature Coasian dynamics. Nevertheless, it is worth noting that in our setting, whether the equilibrium is chosen to minimize or maximize the seller's profit plays a role, as static settings (where some form of the minmax theorem typically holds) do not feature such dramatic discontinuities (see Brooks and Du (2020)).[11]

6. CONCLUSION

This paper has grappled with the possibility of a dynamically-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 benchmark 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 driving equilibria; 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 reinforcing 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. At the same time, our analysis of the no-gap case suggests there are significant subtleties involved with maintaining a unique outcome under the limited commitment approach. The lack of such uniqueness means that the analogy to the known-values case, which necessarily emerges under commitment, requires added assumptions in order to deliver. A natural question is what kinds of refinements might eliminate the possible multiplicity issues and derive intuitive solutions. Compelling solutions would help bridge the significant gap between the robust and Bayesian approaches to mechanism design in the ability to speak to questions of limited commitment.

[11]Note that, since information is specified to depend on the price in the single-period model, the outcome does not depend on if the seller moves first or nature moves first, provided this "richer" action space for nature is still allowed. Without this added richness, randomization may be necessary.

References

Lawrence Ausubel, Peter Cramton, and Peter Deneckere. 2002. Bargaining with incomplete information. In *Handbook of Game Theory with Economic Applications*. Vol. 3. Elsevier, 1897–1945.

Lawrence M. Ausubel and Raymond Deneckere. 1989. Reputation in Bargaining and Durable Goods Monopoly. *Econometrica* 57, 3 (1989), 511–531.

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.

Dirk Bergemann and Stephen Morris. 2005. Robust Mechanism Design. *Econometrica* 73 (2005), 1771–1813.

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

Lukas Bolte and Gabriel Carroll. 2020. Robust contracting under double moral hazard. *Working Paper, Stanford University and University of Toronto* (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. 2020. A Strong Minimax Theorem for Informationally-Robust Auction Design. *Working Paper, University of Chicago and University of California, San Diego* (2020).

Benjamin Brooks and Songzi Du. 2021. Optimal Auction Design with Common Values: An Informationally Robust Approach. *Econometrica* 89, 3 (2021), 1313–1360.

Gabriel Carroll. 2015. Robustness and linear contracts. *American Economic Review* 105, 2 (2015), 536–563.

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

Kim-Sau Chung and Jeffrey Ely. 2007. Foundations of dominant-strategy mechanisms. *Review of Economic Studies* 74 (2007), 447–476.

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).

Larry Epstein and Martin Schneider. 2007. Learning under ambiguity. *Review of Economic Studies* 74, 4 (2007), 1275–1303.

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.

Niccolo Lomys. 2018. Learning while Bargaining: Experimentation and Coasean Dynamics. *Working Paper, Tolouse 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).

Sergio Ocampo Diaz and Keler Marku. 2019. Robust contracts in common agency games. *Working Paper, Western University and University of Minnesota* (2019).

Alessandro Pavan. 2017. Dynamic Mechanism Design: Robustness and Endogenous Types. In *Advances in Economics and Econometrics*. Vol. 1. Cambridge University Press, 1–62.

Doron Ravid, Anne-Katrin Roesler, and Balazs 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* (2020).

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.

A. PROOFS FOR SECTION 3

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 $s = (s_1, \ldots, s_{T-1})$ 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.[12] In this case it is without loss to

[12]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.

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 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. □

Proof of Proposition 1. Suppose the seller chooses price p_1 in period 1, and suppose a candidate equilibrium required nature to choose information structure $\mathcal{I}_1$. To deter a deviation from nature, we suppose that the continuation play following nature's deviation is as good as possible from the seller. Therefore, if $(\hat{p}_2, \hat{p}_3, \ldots,)$ is the conjectured price path the buyer would imagine the seller would use, following this deviation, then we can define $\hat{w}(p_1)$ to be the expected value of the buyer which would be indifferent between buying and not, assuming no further information:

$$\hat{w}(p_1) - p_1 = \max_{\tau} \delta^{\tau}(\hat{w}(p_1) - \hat{p}_{\tau})$$

Note that if the buyer were to receive information in future periods, then this would make delay more attractive, therefore making the buyer *strictly* prefer delay to purchase. On the other hand, since $p_1 > \delta \mathbb{E}[p_2]$ in any equilibrium (as argued in the Proof of Theorem 1).

So, let $w^*(p_1) = \inf_{(\hat{p}_2, \hat{p}_3, \ldots)} \hat{w}(p_1)$. Consider deviations of nature from $\mathcal{I}_1$ where the buyer is told whether v is above or below $F^{-1}(G(w^*(p_1))) - \varepsilon$, for $\varepsilon \to 0$. Then if v is below this threshold, there is no conjecture the buyer could make about the seller's future behavior which would lead them to want to purchase, by the definition of $w^*(p_1)$. On the other hand, above this threshold, for ε small, we will have $\mathbb{E}[v \mid v \geq F^{-1}(G(w^*(p_1))) - \varepsilon] > w^*(p_1)$ and therefore the buyer will buy, given this conjecture.

So, by deviating in this way, nature has the ability to enusre the seller only obtains $p_1(1 - G(w^*(p_1)))$ in period 1. With this in mind, let $\overline{v}_t$ be the highest consumer value that has not purchased by time t, and $\overline{y}(p_1)$ thecorresponding choice of nature. We then have the following recursive formulation for an upper bound of the seller's profit, for every p_1:

$$V(\overline{v}_t) = p_1(1 - G(w^*(p_1))) + \delta V(\overline{y}(p_1))$$

This is precisely the value function in the known values case of Theorem 4 of Ausubel et al. (2002), when the buyer's value is distributed according to G. While we emphasize that the above expression has a less direct interpretation—namely, as an upper bound on the equilibrium profit—nevertheless the result immediately implies that the equilibrium values of V_t must be equal to $\overline{v}$ at some t.

Therefore, the following pair of claims will deliver the proposition:

Claim 1: If $V_t = \underline{v}$, then the market is cleared in finite time.

Proof of Claim 1: Note that there exists a range of choices of p for the seller, say $[\underline{v}, p^*]$, such that the seller optimally clears the market in the next period after charging a price of p. Indeed, since nature can always choose the threshold $F^{-1}(G(p))$, then there exists a range where, if the seller were to charge a price in this range, nature could ensure this price lead to a total payoff of less than $\underline{v}$ (since this holds under known values). Therefore, were we to have an equilibrium with $V_t = \underline{v}$ for infinitely many periods, then we must also have $p_s \to p^{**}$ for some $p^{**} \geq p^*$, and the seller obtaining a total discounted payoff of p_s under the information structure. On the other hand, for any distribution of expected values of the buyer, the previous proof shows that if the seller chooses the Coasian price path, the worst-case the seller can obtain is strictly larger than $\underline{v}$. Therefore, the seller would have a profitable deviation in any such equilibrium.

Claim 2: If F satisfies the Lipschitz condition of Theorem 4 of Ausubel et al. (2002), then so does G.

Proof of Claim 2: Note that, for every quantile q, the pressed distribution satisfies $G^{-1}(1-q) < F^{-1}(1-q)$, but the bottom of the support is the same for each. Therefore, we have $G^{-1}(1-q) - \underline{v} < F^{-1}(1-q) - \underline{v}$, so that if F satisfies the Lipschitz condition—i.e., $F^{-1}(q) - \underline{v} \leq Lq$ for all q and some L—then so does G.[13]

Finishing the proof: Claim 2 shows that the upper bound derived above does indeed ensure that $V_t = \underline{v}$ in finite time, since this result holds given any distribution under known values satisfying the Lipschitz condition. Claim 1 therefore shows that, since $V_t = \underline{v}$ in every equilibrium, and with an upper bound existing on the number of periods this takes, it therefore follows that there exists an upper bound by which the market has cleared. This shows that in the case of a gap, the infinite horizon game coincides with the outcome of a sufficiently long finite horizon game, completing the proof. □

Proof of Proposition 2. We consider two cases for this proof; in the first case we take $T = \infty$ and in the second case we take $T < \infty$. The idea behind the construction in both cases is the following:

- On-path, the seller seller chooses a price equal to the buyer's expected value, and no information is provided.
- Meanwhile, the buyer randomizes purchase so that the seller has incentives to follow the equilibrium strategy.

[13]The Ausubel et al. (2002) is stated slightly differently, namely that $v(q) - v(1) \leq L(1-q)$ for all q and some L. Here, $v(q)$ is a decreasing function, representing the value of the buyer at the $1-q$ quantile (so that $v(1)$ is the value of the buyer at the 0th quantile, i.e., $\underline{v}$). In our notation, $F(v)$ is the probability the buyer's value is below v, so that $F(v) = 1 - q$. Our definition therefore replaces q with $1-q$ and v with F^{-1}.

- If the seller deviates, the equilibrium reverts to the worst-case outcome outlined in Theorem 1.

We emphasize that there is no choice of nature to consider, as this is simply exhibiting some information structure where there is no bound on the market clearing time.

We now walk through the details more precisely. Take the strategy exactly as above. We already know that following a deviation, buyer's strategy forms an equilibrium. Indeed, since the buyer's purchasing decision does not depend on their value, the on-path distribution of v conditional on not having purchased at time t is simply F. Thus, the buyer's problem is completely unchanged relative to the case considered in Theorem 1. The seller's continuation strategy following a deviation also forms an equilibrium, by construction. Note that, since we assume the buyer randomizes, note that it is not possible for them to deviate, since all actions occur with positive probability on-path.

Therefore, letting $\pi^*(G)$ denote the profit achieved in the equilibrium from Theorem 1, the seller obtains at most $\pi^*(G)$ following a deviation. Suppose we seek an equilibrium where the seller's continuation value is v at every point in time, for $v^* > \pi^*(G)$. In this case, the buyer purchases with probability ρ at every point in time, where ρ satisfies:

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

where $\rho \in (0,1)$ whenever $v \in (\pi^*(G), \mathbb{E}_{v\sim F}[v]$

Thus, by charging $\mathbb{E}_F[v]$, the seller obtains a higher payoff than what they could obtain from deviating. We thus verify the conditions are satisfied in the proposition: First, the seller uses a constant price path. Second, the profit obtains is the arbitrary $v^* \in (\pi^*(G), \mathbb{E}_{v\sim F}[v])$. And lastly, the market does not clear by any finite time; since ρ is constant, the probability the buyer has not bought at or before time K is $(1-\rho)^K > 0$.

In the case of a finite horizon, the proof is identical except in the last period, we assume the buyer purchases with probability $v/\mathbb{E}_F[v]$; here, we note that the seller's minmax continuation payoff following a deviation is time dependent, although no matter what the time horizon is it is always strictly bounded away from $\mathbb{E}_F[v]$ (indeed, it is always lower than the seller's static monopoly profit, which is lower than $\mathbb{E}_F[v]$). Accommodating this is straightforward and thus omitted. □

B. 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:

$$\int_{\underline{v}}^{y_t} (v - p_t) f(v) dv = \sum_{s=t+1}^{t^*} \delta^{s-t} \left(\int_{y_s}^{y_{s-1}} (v - p_s) f(v) dv \right). \tag{3}$$

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

$$\sum_{s=1}^{t^*} p_s(F(y_{s-1}) - F(y_s)). \tag{4}$$

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.[14] 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 (4) with respect to y_{t+1}, using (3) to implicitly differentiate $y_t(y_{t+1})$. The derivative of the right hand side of (4) 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 (3) 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:

$$\delta(\overline{v}_{t+1} - y_t)f(y_{t+1}) = (y_t - \overline{v}_t)f(y_t)y_t'(y_{t+1}). \tag{5}$$

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

[14] 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.

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

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 (5) 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:

$$\overline{v}_t(\overline{v}_{t+1} - y_{t+1}) + \overline{v}_{t+1}(y_t - \overline{v}_t) = -\overline{v}_t y_{t+1} + \overline{v}_{t+1} y_t. \tag{6}$$

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 (6) 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 (6) 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 (6) 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 Proposition 3. 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:

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

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. □

Proof of Proposition 4. We first describe the set of equilibria delivering the Ausubel and Deneckere (1989) folk theorem under the pressed distribution G. In fact, for reasons that will become clear in the course of the proof, we will do this assuming the buyer obtains an arbitrary initial signal I_0. Note that in this case, we can define a distribution $\tilde{G}_{I_0}$ via the following:

First, let s denote an arbitrary signal realization under I_0, and let F_s denote the distribution of the buyer's value conditional on observing s, and let G_s denote the pressed version of the distribution F_s. We define

$$\tilde{G}_{I_0}(x) = \mathbb{E}_{s \sim I_0}[G_s(x)].$$

Note that, if the buyer were to observe I_0, then conditional on the signal observed, the worst-case information structure conditional on s would be a partitional threshold at $F_s^{-1}(G_s(p))$. Therefore, $\tilde{G}_{I_0}(p)$ defines a distribution such that the probability of sale in the worst-case information structure following a price of p is $1 - \tilde{G}_{I_0}(p)$, if the buyer were to have I_0 before purchase. Note further that, since nature could always provide the signal I_0, by construction we have that the optimal profit following I_0 is weakly higher than the optimal profit following no information, for any candidate equilibrium path.

In fact, given an arbitrary price path for the seller, $p_1, \ldots, p_n, \ldots$, the value which is indifferent between purchasing and not assuming no further information does not directly depend on the signal observed, since this indifference condition only depends on the price path, the expected value of the buyer, and δ. Therefore, the seller's profit from such a price path coincides with the known-values profit under distribution $\tilde{G}_{I_0}(x)$. In the dynamic threshold information structure, where the buyer's expected value conditional on not purchasing is exactly this indifferent value, as long as the buyer follows the recommendations of nature, we again have the seller's profit is the known values profit.

We now show that the conditions for the folk theorem of Ausubel and Deneckere (1989) hold, meaning that, via the above argument, their specification for the equilibrium price path delivers the same profit under that price path in the known values benchmark where the buyer's value is distributed according to $\tilde{G}_{I_0}(p)$. Their condition stated for the known value case is that there exists L, M such that, for all q:

$$Mq^\alpha \leq F^{-1}(q) \leq Lq^\alpha$$

Our claim will follow from the assumption that this condition holds for the pressed distribution and:

$$G^{-1}(q) \leq \tilde{G}_{I_0}^{-1}(q) \leq F^{-1}(q).$$

In that case, we can ensure that, uniformly over the set of information structures I_0, if M is taken from $G^{-1}(q)$ and L is taken from $F^{-1}(q)$, then $Mq^\alpha \leq \tilde{G}_{I_0}^{-1}(q) \leq Lq^\alpha$. This claim, in turn, immediately follows from the definition of G and $\tilde{G}_{I_0}$ as the solution to the worst-case information structure construction from the static case. Indeed, consider the seller choosing *quantiles* instead of prices, so that the seller's profit in the one period problem, facing distribution $\tilde{F}$, is given by $\tilde{F}^{-1}(q) \cdot (1-q)$. Decreasing the quantile given the price decreases the profit; and since nature always has the option of giving I_0 in addition to the threshold, we therefore have $G^{-1}(q) \leq \tilde{G}_{I_0}^{-1}(q)$. Since nature has the option of giving full information instead of the worst-case thresholds following I_0, we have $\tilde{G}_{I_0}^{-1}(q) \leq F^{-1}(q)$.

For the subsequent part of this proof, we let $\underline{\pi}_\delta$ denote the lowest payoff from the above construction (assuming no initial information to the buyer), and we let $\overline{\pi}_\delta(I_0)$ denote the highest possible payoff given an information structure I_0 from the above construction. We note that $\underline{\pi}_\delta \to 0$ and $\overline{\pi}_\delta(I_0)$ converges to the monopoly profit under the "modified" pressed distribution described above, which is weakly large than the monopoly profit under the pressed distribution.

We now turn to the specification of the equilibrium from proposition 2. Specifically we assume that in every period:

- The seller chooses price $p^* = \mathbb{E}[v]$;
- Nature provides no information;
- The buyer randomizes between purchasing and not with probability ρ.

As in the construction of Proposition 2, we note that there is no possible deviation for the buyer since all purchase times occur with positive probability (and at all of them, the payoff obtained is 0).

We consider deviations of the seller and nature, intuitively moving to the best possible equilibrium for the seller in the case where nature deviates and the worst possible equilibrium for the seller if the seller deviates, under the Ausubel and Deneckere (1989) equilibria from above.

That there is no profitable deviation for the seller is immediate; in this case, the equilibrium *immediately* shifts to one where the monopolist's payoff is no more than $\underline{\pi}_\delta$, which by construction is lower than what the seller obtains on-path.

For nature, note that the *best* case for the seller is that the buyer purchases at price $\mathbb{E}[v]$, since this is an upper bound on the surplus the seller could obtain in any equilibrium. Therefore, a lower bound on the seller's profit is achieved by assuming no buyers purchase in that period. In that case, the seller obtains $\delta\overline{\pi}_\delta(I_0)$, which for sufficiently large δ is large than the on-path payoff (since on path, the seller obtains strictly less than the single-period monopoly profit under G, whereas as $\delta \to 1$, $\delta\overline{\pi}_\delta(I_0)$ converges to this amount). Therefore, nature does not want to deviate from the prescribed equilibrium, either, completing the proof. $\square$