

Dynamic Pricing with Limited Commitment^{*}

Martino Banchio[†] Frank Yang[‡]

December 9, 2021

Abstract

A monopolist wants to sell one item per period to a consumer with evolving and persistent private information. The seller sets a price each period depending on the history so far, but cannot commit to future prices. We show that, regardless of the degree of persistence, any equilibrium under a D1-style refinement gives the seller revenue no higher than what she would get from posting all prices in advance.

Keywords: Dynamic pricing, limited commitment, ratchet effect, dynamic mechanism design.

JEL Codes: C73, D42, D82.

^{*}This paper has been previously circulated under the title “Static Pricing.” We thank Mohammad Akbarpour, Anirudha Balasubramanian, Gabriel Carroll, Daniel Chen, Laura Doval, Matthew Gentzkow, David Kreps, Giacomo Mantegazza, Paul Milgrom, Michael Ostrovsky, David Ritzwoller, Ilya Segal, Jesse Shapiro, Andy Skrzypacz, Takuo Sugaya, Bob Wilson, and Weijie Zhong for helpful comments.

[†]Graduate School of Business, Stanford University. Email: mbanchio@stanford.edu

[‡]Graduate School of Business, Stanford University. Email: shuny@stanford.edu

1 Introduction

As electronic commerce and personalized sales become prevalent, dynamic pricing has received a surge of attention (e.g., [Kehoe et al. 2020](#)).¹ Consider a monopolist (she) selling nondurable goods repeatedly to a consumer (he) with evolving private information about his value. The buyer’s private values are positively correlated across periods. The seller adjusts prices dynamically based on the information learned through purchase histories, but cannot commit to future prices. Can the seller improve upon simply posting all prices in advance?

In this paper, we study dynamic pricing with evolving private information and limited commitment. In the benchmark model, we consider a two-period game and show that regardless of the degree of persistence of the private information, any equilibrium satisfying a D1-style refinement gives the seller revenue no higher than what she would get from posting all prices in advance.

This result contrasts sharply with the case of full commitment. A large and growing literature on dynamic mechanism design highlights the value in the ability to dynamically respond to the buyer’s evolving private information ([Baron and Besanko 1984](#), [Courty and Li 2000](#), [Eso and Szentes 2007](#), [Pavan et al. 2014](#)). For example, in a repeated-sales model, [Battaglini \(2005\)](#) shows that for any imperfect correlation of the buyer’s private values, the seller’s average profits converge to the first-best through long-term contracting as the discount factor converges to one.

Instead of analyzing general mechanisms, we restrict attention to (potentially randomized) pricing strategies.^{2,3} By analyzing the equilibria of a dynamic pricing game with evolving private information, we show that the benefits of the flexibility in dynamic pricing hinge critically on the long-term commitment power.

An important special case arises when the consumer’s value distribution is stationary (i.e., identical marginals) but exhibits positive serial correlation. Our result implies that the monopolist cannot do better than simply maintaining the static monopoly price. An extreme case of such a stationary setting is perfect correlation (i.e., constant types). In this extreme case, our result follows from the classic analysis of [Stokey \(1979\)](#), who shows that

¹The Wall Street Journal reports, “The Journal identified several companies, including Staples, Discover Financial Services, Rosetta Stone Inc. and Home Depot Inc., that were consistently adjusting prices and displaying different product offers based on a range of characteristics that could be discovered about the user.” *Websites Vary Prices, Deals Based on Users’ Information*, The Wall Street Journal, December 24, 2012.

²Among all selling formats, pricing is the most common one in online retail markets; for example, roughly 90% of the listings in 2015 on eBay are posted price listings ([Einav et al., 2018](#)).

³In a related but different setting (a single durable good with constant values), [Skreta \(2006\)](#) shows that it is without loss of generality to consider pricing strategies among all sequentially rational mechanisms.

the optimal long-term contract with full commitment is static. However, as pointed out earlier, as soon as the correlation is not perfect, the seller can extract more revenue using long-term contracts by conditioning future prices on the buyer’s previous decisions.

Why does dynamic pricing, so effective when the seller has long-term commitment power, become futile when the seller has limited commitment? With evolving private information, an important lesson from dynamic mechanism design is that early contracting at the initial stage prevents the buyer from capitalizing on his future information rents. However, such early contracting requires substantial commitment power. Because the types are persistent over time, it is sequentially rational for the seller to post a high price following an acceptance of the first-period offer and a low price following a rejection. As in the constant-type case, the *ratchet effect* appears — the buyer anticipates the effect of his present choice on future offers: he could give up the purchase today in exchange for a better price tomorrow. As our analysis shows, if the seller could commit to history-contingent prices, then the price schedule would never be time-consistent.

However, the intuition is incomplete. In environments where more information arrives in the future, the buyer also worries less about his decisions revealing the current private information. The ratchet effect becomes weaker as the persistence of private information decreases. For example, [Kennan \(2001\)](#) studies a version of our bargaining game with binary types. In contrast to the constant-type case of [Hart and Tirole \(1988\)](#), he finds that “when the valuation is persistent, but not fully permanent, the ratchet effect is not so potent: learning does occur in equilibrium.” In light of his analysis, our result shows that with persistent private information, even though learning is possible in equilibrium, it is unprofitable. The ratchet effect is potent enough to eliminate the potential gains from dynamic pricing.

The key trade-off of commitment versus learning requires a precise understanding of how equilibrium payoffs depend on the joint distribution of buyer’s values. However, as shown in [Kennan \(2001\)](#), there is a plethora of equilibria even when the buyer’s types are binary. Characterizing all equilibria in games involving the ratchet effect is often intractable (see e.g., [Laffont and Tirole 1988](#)). Our proof strategy connects the dynamic pricing game back to a constrained mechanism design problem. Instead of characterizing all equilibria, we identify a set of necessary conditions on the allocation rule satisfied in all equilibria, and introduce an auxiliary dynamic mechanism design problem constrained by these necessary conditions. The value of this auxiliary problem provides an upper bound on the seller’s payoff. We show that posting the monopoly price for each good precisely attains the upper bound.

In the benchmark model, we assume that the buyer has additive values across peri-

ods. In [Section 4.1](#), we also show that the same result holds even if the products may be complements. We also provide examples to show that when the values are negatively correlated or when the goods are substitutes, the ratchet effect need not appear, which can then lead to dynamic pricing being profitable even under limited commitment (see [Remark 6](#) and [Remark 8](#)). In [Section 4.2](#), we introduce a multiple-period version of our game in which the seller always has an option to commit to a price for each unsold good, and show that a similar conclusion holds for arbitrary finite time horizons, under a stronger distributional assumption.

While our result is a theoretical demonstration of the detrimental power of the ratchet effect, it also has practical relevance. First, not all sellers in reality have long-term commitment power. Even though dynamic pricing policies such as loyalty cards and reward programs are used in practice, attempts to renege, modify, or even cancel them are not rare. In recent years, several class-action law suits have been filed against loyalty programs of multiple companies, including AutoZone (Hughes et al. v. AutoZone Parts et al., 2016), Staples (Torczyner v. Staples, 2016), and AT&T (Palumbo v. AT&T, 2021). Second, because of these issues, such long-term pricing requires substantial trust which also depends on the cultural values in specific markets (e.g., [Thompson and Chmura 2015](#) suggest that in countries with high uncertainty avoidance culture, consumers tend to find loyalty programs less appealing). Third, machine learning based pricing algorithms, increasingly popular in online market places, often directly incorporate purchase histories as one of the attributes (e.g., [Gupta and Pathak 2014](#)). In all these scenarios, our result suggests that the seller may be better off by posting the price for each good in advance.⁴ Of course, maintaining these prices also requires some commitment power. But arguably such a strategy is more convenient and has fewer trust issues, especially in the important case of stationary value distributions where the seller only needs to maintain a single price (e.g., [Einav et al. 2018](#)).⁵

Our model builds on a large literature on behavior-based price discrimination (e.g., [Hart and Tirole 1988](#), [Villas-Boas 2004](#), [Acquisti and Varian 2005](#)), which is surveyed by [Fudenberg and Villas-Boas \(2006, 2012\)](#) and [Acquisti et al. \(2016\)](#). Most work in this literature assumes that the consumer's values are perfectly correlated across periods.⁶

⁴We also show that even if the seller is allowed to commit to history-dependent prices, as long as she cannot use policies like repeated-purchase discounts by setting the price after acceptance lower than the price after rejection, the optimal mechanism is still to simply post all prices in advance (see [Remark 4](#)).

⁵[Einav et al. \(2018\)](#) document the decline of auctions used in online marketplaces and argue that the convenience of posted prices is one of the main drivers. Posting all prices in advance is also known as price commitment, which has been argued to be a common pricing strategy (see e.g., [Jing 2011](#)).

⁶Two notable exceptions are [Taylor \(2004\)](#) and [Kennan \(2001\)](#). Both papers study a binary-type setting and allow for positive correlation of types. [Taylor \(2004\)](#) characterizes the pure-strategy equilibria; [Kennan](#)

This assumption rules out both the possibility of differentiated goods and the possibility of taste shocks. More importantly, as indicated earlier, the case of perfect correlation is a knife-edge case in which the effectiveness of behavior-based price discrimination is muted even with full commitment. Our result shows that the detrimental power of the ratchet effect, however, is not knife-edge — it eliminates the gains from dynamic pricing for any degree of persistence even when the optimal contract itself is dynamic.

The study of the ratchet effect goes back to the classic work of [Freixas et al. \(1985\)](#), [Laffont and Tirole \(1988\)](#), and [Hart and Tirole \(1988\)](#). Recently, there has been a renewed interest in better understanding its impact (e.g., [Gerardi and Maestri 2020](#), [Bonatti and Cisternas 2020](#)). In particular, [Bonatti and Cisternas \(2020\)](#) study how strategic consumers interacting with a sequence of firms (in continuous time) reduce their demand to manipulate their scores that aggregate past purchase histories. In contrast, in our model, the monopolist is long-lived and can condition her price offer on the entire sequence of purchase histories.

Our result connects the ratchet effect to the growing literature of dynamic mechanism design, which is surveyed by [Pavan \(2017\)](#) and [Bergemann and Välimäki \(2019\)](#). As [Pavan \(2017\)](#) points out, “[The limited commitment] 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.” Indeed, as our result shows, introducing limited commitment can give a qualitatively different answer to important economic questions such as whether dynamic pricing is profitable. While we restrict attention to randomized pricing mechanisms (instead of general mechanisms), this approach enables us to (i) provide an empirically relevant benchmark result and (ii) connect the economics of the ratchet effect, which is at the heart of limited commitment models, to dynamic mechanism design.^{7,8} We also show that with evolving private information, equilibrium refinement plays a crucial role in modeling limited commitment (see [Remark 2](#)), and introduce a notion of D1 criterion in the spirit of [Banks and Sobel \(1987\)](#) for our game (see [Appendix B](#)).

(2001) shows that there is often plethora of mixed-strategy equilibria and characterizes a particular class of them. In contrast, our model allows for a continuum of types in each period, and our result holds for all mixed-strategy equilibria under a D1-style refinement.

⁷Our approach is conceptually similar to [Liu et al. \(2019\)](#) who study limited commitment while restricting attention to auctions. In contrast to us, they study a durable good monopoly setting in which the game ends after the good is traded, and do not introduce evolving private information.

⁸A different form of limited commitment is considered by [Deb and Said \(2015\)](#) who focus on the arrival of new consumers instead of the evolution of private information. Also, see [Doval and Skreta \(2020\)](#) for general mechanism design with limited commitment. Our approach has the additional benefit of being tractable. Unlike [Strulovici \(2017\)](#), [Doval and Skreta \(2019\)](#), and [Gerardi and Maestri \(2020\)](#), our model allows for a continuum of types in each period.

Our paper also relates to a more applied literature on dynamic pricing and price commitment, especially for experience goods for which the consumer learns his value after consuming the first unit (e.g., [Cr mer 1984](#) and [Jing 2011](#)). The key difference is that in such models, the consumer is often ex ante uninformed, and the purchase history only reveals whether the consumer becomes informed about his value. Thus, the ratchet effect is absent. Our model allows for consumption-independent learning and focuses on the ratchet effect induced by the positive correlation of values. Of course, in doing so, we abstract away many other important features of repeated sales (e.g., finite inventories, stochastic arrivals, and search costs). There is a large literature on dynamic pricing in management science that focuses on these features (e.g., [Elmaghraby and Keskinocak 2003](#), [Su 2007](#), [Aviv and Pazgal 2008](#), [Cachon and Feldman 2015](#)).

The rest of the paper proceeds as follows. [Section 2](#) presents the model. [Section 3](#) presents the main result and its proof. [Section 4](#) shows the robustness of the result to two variations: complementary products ([Section 4.1](#)) and multiple periods ([Section 4.2](#)). [Section 5](#) concludes.

2 Model

A seller wants to sell a nondurable good in each period $t \in \{1, 2\}$ to a buyer, with a cost normalized to 0. The buyer privately observes his value $\theta_t \in [\underline{\theta}_t, \bar{\theta}_t] \subset \mathbb{R}_+$ at the beginning of each period. Let $F_1(\theta_1), F_2(\theta_2)$ denote the marginal distributions of θ_1, θ_2 respectively, and $F_2(\theta_2|\theta_1)$ denote the distribution of θ_2 conditional on θ_1 . The distributions are commonly known. We assume that $F_1(\theta_1)$ is absolutely continuous with a density denoted by $f_1(\theta_1)$, and $F_2(\theta_2|\theta_1)$ is continuously differentiable in (θ_1, θ_2) with a conditional density denoted by $f_2(\theta_2|\theta_1)$. We also assume that $f_1(\theta_1) > 0$ for all $\theta_1 \in [\underline{\theta}_1, \bar{\theta}_1]$.

At the beginning of each period t , the seller makes a take-it-or-leave-it offer p_t to the buyer for the product in period t . The buyer decides whether to accept the offer. After the buyer makes his decision, the game moves to the next period. Let $x_t \in \{0, 1\}$ denote the decision of the buyer in period t . We assume that both the seller and buyer have quasilinear preferences that are additively separable across the two periods, with a common discount factor $\delta \in (0, 1]$. So the seller's period- t payoff is given by $u_t^S(x_t, p_t) = \delta^{t-1} x_t p_t$, and the buyer's period- t payoff is given by $u_t^B(\theta_t, x_t, p_t) = \delta^{t-1} x_t (\theta_t - p_t)$.

Strategies and Beliefs

Let h_t and \hat{h}_t denote the public history observed by the seller and the history observed by the buyer up to period t , respectively. The seller's strategy σ maps a public history h_t to a distribution on \mathbb{R} representing the price announced in period t . A buyer's strategy τ maps a history \hat{h}_t to a distribution on $\{0, 1\}$ representing his purchase decision in period t . The seller forms belief $\mu(h_t)$ about $\theta := (\theta_1, \theta_2)$ given a public history h_t .

Equilibrium

Our choice of equilibrium concept is perfect Bayesian equilibrium with a further restriction on the seller's off-path beliefs. A perfect Bayesian equilibrium satisfying the D1 criterion (PBE-D) of this game is a tuple (σ, τ, μ) such that the following hold:

- Given $\mu(h_t)$ and τ , σ is sequentially rational at every public history h_t .
- Given σ , τ is sequentially rational at every history \hat{h}_t .
- $\mu(h_t)$ is derived via Bayes' rule whenever possible.
- Seller's off-path beliefs satisfy a D1 criterion, formally defined in [Appendix B](#).

The only difference from a standard PBE is in the last point. We require that an equilibrium satisfy a D1 criterion in the spirit of [Banks and Sobel \(1987\)](#), at any off-path history. The seller cannot update her belief via Bayes' rule at some history h_t when all types of the buyer accept (or reject) in equilibrium, but she observes a rejection (or acceptance). Intuitively, the refinement requires that when observing an unexpected deviation, the seller believes that the buyer is of the type that could benefit the most from the deviation. Even though our model is not a signaling game, [Banks and Sobel \(1987\)](#)'s notion can be extended to our setting, as formally shown in [Appendix B](#).

Distributional Assumptions

We impose the following assumptions for our main result.

Assumption A1 (MLRP). For any $\theta'_1 > \theta_1$, the conditional likelihood ratio $\frac{f_2(\theta_2|\theta'_1)}{f_2(\theta_2|\theta_1)}$ is strictly increasing in θ_2 .

[Assumption A1](#) is equivalent to strict affiliation of (θ_1, θ_2) as in [Milgrom and Weber \(1982\)](#). For example, consider an AR(1) process of the form $\theta_2 = \alpha\theta_1 + \epsilon$, where ϵ is drawn from a Gaussian distribution. Then [Assumption A1](#) is equivalent to $\alpha > 0$.

Assumption A2 ($1/\delta$ -Lipschitz). For any $\theta'_1 > \theta_1$,

$$\mathbb{E}[\theta_2|\theta'_1] - \mathbb{E}[\theta_2|\theta_1] < \frac{1}{\delta}(\theta'_1 - \theta_1).$$

Assumption A2 says that a change in the first-period type leads to a comparable change in the second-period type in expectation. For any AR(1) process $\theta_2 = \alpha\theta_1 + \epsilon$, **Assumption A2** is equivalent to $\alpha < \frac{1}{\delta}$.

To state our last assumption, we introduce the notion of an impulse response function (Pavan et al. 2014):

$$I(\theta_1, \theta_2) := -\frac{\partial F_2(\theta_2|\theta_1)/\partial \theta_1}{f_2(\theta_2|\theta_1)}.$$

Intuitively, the impulse response function measures the effect of a small change in θ_1 on θ_2 . For any AR(1) process, $\theta_2 = \alpha\theta_1 + \epsilon$, we have $I(\theta_1, \theta_2) = \alpha$.

Assumption A3 (Regularity). The second-period virtual value function

$$\psi(\theta_1, \theta_2) := \theta_2 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I(\theta_1, \theta_2)$$

is non-decreasing in both θ_1 and θ_2 .

Assumption A3 is the dynamic version of Myerson's regularity condition and commonly imposed in dynamic mechanism design for tractability (see e.g., Bergemann and Välimäki 2019). For any AR(1) process, we have $\psi(\theta_1, \theta_2) = \theta_2 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)}\alpha$, and hence **Assumption A3** holds if the distribution F_1 has monotone hazard rate.

Remark 1. Under the distributional assumptions, we show in **Appendix B** that the D1 criterion is satisfied if the seller believes the buyer is of type $\underline{\theta}_1$ when observing an off-path rejection, and of type $\bar{\theta}_1$ when observing an off-path acceptance. Moreover, we also show in **Appendix B** that for any PBE that satisfies the D1 criterion, there exists an outcome-equivalent PBE with off-path beliefs of this form.⁹ Therefore, for all of our results, it is without loss of generality to focus on the PBE with such off-path beliefs. See **Remark 2** for the importance of restricting off-path beliefs in modeling limited commitment.

3 Main Result

Let $\pi_t^* = \max_p p(1 - F_t(p))$ be the standard monopoly revenue with respect to the marginal distribution in period t . Our main result says that any equilibrium gives the seller rev-

⁹We say two equilibria are outcome-equivalent if they almost surely induce the same allocation and the same total discounted payment.

enue no higher than what she would get by simply posting all prices in advance:

Theorem 1. Under *Assumptions A1 to A3*,

(i) *there exists a PBE-D;*

(ii) *in every PBE-D, the seller's revenue is no higher than $\sum_t \delta^{t-1} \pi_t^*$.*

We leave the proof of part (i) of [Theorem 1](#) to [Appendix A.6](#). The proof proceeds by considering an auxiliary one-shot continuous game, showing the existence of an equilibrium in the auxiliary game using the result of [Glicksberg \(1952\)](#), and constructing a PBE-D using the continuity properties of the equilibria of the auxiliary game.

We prove part (ii) of [Theorem 1](#) in [Section 3.1](#). The intuition behind the proof of this revenue bound can be understood as follows. Suppose hypothetically that the seller could commit to any history-contingent price path that specifies the first-period price p_1 , the second-period price after an acceptance p_A , and the second-period price after a rejection p_R . We show that the seller would always commit to some $p_A < p_R$. Before explaining this claim, we first explain why this implies that the seller's revenue in equilibrium is no higher than the sum of discounted monopoly revenues.

When the seller observes an acceptance, she believes that the buyer has a high type. Because the types are persistent, sequential rationality implies that the seller always sets $p_A \geq p_R$ in equilibrium. But by our earlier claim, the seller would prefer to set $p_A < p_R$. Thus, the seller with limited commitment cannot obtain a revenue higher than what she would get by letting $p_A = p_R$, or equivalently by simply posting all prices in advance. [Figure 1](#) illustrates this argument with stationary distributions. When types are perfectly correlated, the optimal commitment prices p_A, p_R are both equal to the static monopoly price p^* . When the correlation becomes imperfect, the seller benefits from committing to a lower p_A and a higher p_R , but sequential rationality constrains the seller in exactly the *opposite* direction.

To see why the seller would prefer to set $p_A < p_R$ if allowed to commit, it is helpful to consider the following simple example (assuming $\delta = 1$):

Example 1. The buyer has independent values across the two periods, and in each period θ_t is drawn uniformly from the interval $[1, 2]$.

Note that the monopoly price for the uniform distribution on $[1, 2]$ is 1. Thus, if the seller posts prices in advance, then she will post price 1 in both periods and collect a total revenue of 2. Now, suppose the seller commits to $(p_1 = 1.5, p_A = 1, p_R = 2)$. Consider how the buyer of type θ_1 decides in the first period. If he rejects the first period offer, then he

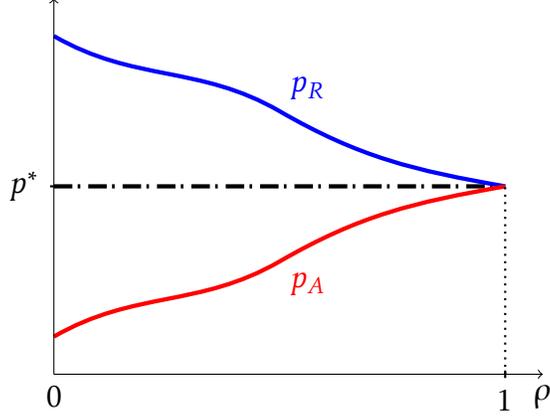


Figure 1: Illustration of the seller's optimal commitment prices

gets 0 payoff because the second-period offer is too high. If he accepts, then his expected payoff is given by

$$\theta_1 - p_1 + \mathbb{E}[(\theta_2 - p_A)_+ | \theta_1] = \theta_1 - p_1 + \mathbb{E}[\theta_2 - p_A | \theta_1] = \theta_1 - 1.5 + \mathbb{E}[\theta_2] - 1 = \theta_1 - 1 \geq 0$$

because θ_2 is always no less than $p_A = 1$ and θ_2 is independent of θ_1 . Therefore, all types of buyer will accept the first-period offer and then accept the second-period offer. The seller gets a total revenue of $1.5 + 1 = 2.5 > 2$.

The first-period price $p_1 = 1.5$ includes both the monopoly price for the first good and an option price of 0.5 for getting the second good at $p_A = 1$. By committing to a low p_A and a high p_R , the seller can add a surcharge on top of the monopoly price to sell the second good “in expectation.” The buyer only gets information rents from his private information about θ_1 , as the second-period surplus $\mathbb{E}[\theta_2] = 1.5$ is completely extracted.

This observation does not hinge on the assumption that the types are independent across periods. As long as there is some uncertainty in the second-period types, the seller in general benefits from committing to some $p_A < p_R$. This effect echoes the power of “early contracting” in the literature on dynamic mechanism design: contracting before the agent's future private information realizes allows the designer to extract more surplus. In the proof, we use the techniques from dynamic mechanism design to make a precise statement about when the seller prefers to commit to $p_A < p_R$.

Remark 2. The sequential rationality constraint only comes into play if the seller is at an on-path history. Without any equilibrium refinement, the seller can have implausible beliefs at some off-path history in equilibrium. Such beliefs can loosen the sequential rationality constraint and effectively grant the seller more and sometimes full commitment power. To see this concretely, consider the following example:

Example 2. Type θ_1 is drawn uniformly from $[1, 2]$. If type $\theta_1 = 2$, then type $\theta_2 = 2$; otherwise, θ_2 is drawn uniformly from $[1, 2]$ independently of θ_1 .

[Example 2](#) differs from [Example 1](#) only when $\theta_1 = 2$, but the seller can now obtain the full-commitment outcome that we have just described ($p_1 = 1.5, p_A = 1, p_R = 2$) in a perfect Bayesian equilibrium if there is no restriction on the off-path beliefs.¹⁰ All types solve the same problem in the first period as before, except the type $\theta_1 = 2$, but this type also prefers accepting the first-period offer. So all types accept in both periods. Now suppose the seller holds a belief μ concentrated on $\theta_1 = 2$ after a first-period rejection (which is an off-path history). This belief makes the price $p_R = 2$ sequentially rational, effectively granting the seller full commitment power. The D1 refinement eliminates such implausible beliefs by essentially requiring the seller to believe that the buyer who deviates must be of the type most likely to profit from the off-path deviation. As noted in [Remark 1](#), the seller would believe that the buyer is of type $\theta_1 = 1$ after observing an off-path rejection, resulting in $p_A \geq p_R$ even at an off-path history.

3.1 Proof of the Main Result

The proof of the revenue bound, part (ii) of [Theorem 1](#), proceeds in three steps: (1) identify a set of necessary conditions that are satisfied in every PBE-D, (2) consider a mechanism design problem constrained by this set of necessary conditions, and (3) show that the optimal solution to this relaxed problem can be implemented by simply posting all prices in advance. For expositional convenience, in this section we prove the statement for all pure-strategy PBE-D. In [Appendix A.3](#), we extend the proof to cover all PBE-D.

Step 1: Let p_A, p_R be the seller's equilibrium choice of prices posted after observing an acceptance and a rejection, respectively. Let $x_1(\theta_1)$ be the equilibrium first-period allocation. We show that a key restriction imposed by sequential rationality is the following set of monotonicity constraints:

Proposition 1. *In every pure-strategy PBE-D, x_1 is a threshold rule and $p_A \geq p_R$.*

Before proving this claim, we collect two lemmas.

¹⁰This outcome is actually optimal even if the seller could commit to any mechanism. To see this, note that the seller can never do better than observing the realization of θ_2 and then committing to a direct mechanism mapping a type report $\hat{\theta}_1$ to an allocation and a transfer. In that case, the optimal mechanism is to allocate both goods and charge a price of $1 + \theta_2$, which generates a revenue of 2.5.

Lemma 1. For any $k \in (\underline{\theta}_1, \bar{\theta}_1)$,

$$\frac{f(\theta_2|\theta_1 \geq k)}{f(\theta_2|\theta_1 < k)}$$

is strictly increasing in θ_2 , i.e., $F_2(\cdot|\theta_1 \geq k) \succ_L F_2(\cdot|\theta_1 < k)$.¹¹

Proof. See [Appendix A.1](#). □

This lemma is a direct consequence of [Assumption A1](#). It sorts the seller's posterior belief in the second period, assuming a threshold rule in the first period.

Lemma 2. If $F_2 \succ_L F_1$ and p_i is an optimal monopoly price under F_i , then $p_2 \geq p_1$.

Proof. See [Appendix A.2](#). □

Because likelihood ratio ordering implies inverse hazard rate ordering, the above observation follows immediately from the monotone selection theorem of [Topkis \(1998\)](#).

Proof of Proposition 1. In the second period, the buyer accepts when the price is below his second-period type θ_2 . Thus, in the first period, type θ_1 accepts a price p only if

$$(\theta_1 - p) + \delta \mathbb{E}[(\theta_2 - p_A)_+|\theta_1] \geq \delta \mathbb{E}[(\theta_2 - p_R)_+|\theta_1]. \quad (1)$$

Let $m(\theta_2) := (\theta_2 - p_R)_+ - (\theta_2 - p_A)_+$, $h(\theta_1) := \mathbb{E}[m(\theta_2)|\theta_1]$ and $g(\theta_1) := \frac{1}{\delta}(\theta_1 - p)$. So type θ_1 accepts only if $g(\theta_1) \geq h(\theta_1)$ and always accepts if the inequality is strict. Note that $m(\theta_2)$ is 1-Lipschitz and has a.e. derivative $m' \leq 1$. For any $\theta'_1 > \theta_1$, integration by parts yields

$$\begin{aligned} h(\theta'_1) - h(\theta_1) &= \mathbb{E}[m(\theta_2)|\theta'_1] - \mathbb{E}[m(\theta_2)|\theta_1] = \int_0^\infty [\mathbb{P}(\theta_2 > s|\theta'_1) - \mathbb{P}(\theta_2 > s|\theta_1)] m'(s) ds \quad (2) \\ &\leq \mathbb{E}[\theta_2|\theta'_1] - \mathbb{E}[\theta_2|\theta_1] < \frac{1}{\delta}(\theta'_1 - \theta_1) = g(\theta'_1) - g(\theta_1) \end{aligned}$$

where we have also used $\theta_2|\theta'_1 \succeq_{FOSD} \theta_2|\theta_1$ (implied by [Assumption A1](#)) and the Lipschitz condition on $\mathbb{E}[\theta_2|\cdot]$ ([Assumption A2](#)). Hence, g crosses h at most once from below. The first-period allocation is characterized by a cutoff type k , proving the first part.

Suppose the cutoff type $k \in (\underline{\theta}_1, \bar{\theta}_1)$. By [Lemma 1](#), $F_2(\cdot|\theta_1 \geq k) \succ_L F_2(\cdot|\theta_1 < k)$. Then by [Lemma 2](#), $p_A \geq p_R$, because p_A must be optimal under belief $F_2(\cdot|\theta_1 \geq k)$ and p_R must be optimal under belief $F_2(\cdot|\theta_1 < k)$. If $k = \underline{\theta}_1$, then a rejection is off-path and it is without loss of generality to let the seller form a belief $F_2(\cdot|\theta_1 = \underline{\theta}_1)$ as noted in [Remark 1](#). Hence, $p_A \geq p_R$ as $F_2(\cdot|\theta_1 \geq \underline{\theta}_1) \succ_L F_2(\cdot|\theta_1 = \underline{\theta}_1)$. The same argument holds for $k = \bar{\theta}_1$. □

¹¹We use the notation \succ_L for strict likelihood-ratio dominance, i.e., the likelihood ratio of two distributions is a strictly increasing function.

Step 2: Because the two monotonicity constraints must hold in any equilibrium, if the seller can commit to any price mechanism subject to these constraints, then she can always replicate any (pure-strategy) PBE-D. Therefore, providing her the commitment power subject to these constraints forms a relaxed problem.

Specifically, the relaxation is done as follows:

Original:	Relaxed:
Select a seller-optimal PBE-D	Design a revenue-maximizing mechanism
1. Buyer's IC constraints	1. Buyer's IC constraints
2. Seller's sequential rationality	2. $x_1(\theta_1) = \mathbb{1}_{\theta_1 \geq k}$
3. Belief consistency	3. $x_2(\theta_1, \theta_2) = x_1(\theta_1)\mathbb{1}_{\theta_2 \geq p_A} + (1 - x_1(\theta_1))\mathbb{1}_{\theta_2 \geq p_R}$, $p_A \geq p_R$

In this relaxation, we replace the sequential rationality and belief consistency constraints with the constraints imposed directly on the allocation rules. The space of price mechanisms also constrains the set of feasible allocation and payment rules jointly. Let \mathcal{X} denote the set of all incentive-compatible dynamic direct mechanisms satisfying the feasibility and monotonicity constraints. The seller then solves the following problem:

$$\max_{(x,t) \in \mathcal{X}} \mathbb{E} \left[\theta_1 x_1(\theta_1) + \delta \theta_2 x_2(\theta_1, \theta_2) - V_{x,t}(\theta_1) \right] \quad (3)$$

where $V_{x,t}(\theta_1)$ is the expected payoff of a type θ_1 buyer under the mechanism (x, t) . (We use t to denote both the transfer and time. Its meaning should be clear from the context.) As in static mechanism design, local incentive compatibility constraints imply a version of payoff equivalence in dynamic environments. In particular, we use the following result from [Pavan et al. \(2014\)](#):

Proposition 2 ([Pavan et al. 2014](#)). *For any incentive-compatible direct mechanism (x, t) , $V_{x,t}(\theta_1)$ is Lipschitz continuous with the derivative given almost everywhere by*

$$V'_{x,t}(\theta_1) = \mathbb{E} \left[x_1(\theta_1) + \delta I(\theta_1, \theta_2) x_2(\theta_1, \theta_2) \middle| \theta_1 \right]. \quad (4)$$

Proof. See Theorem 1 in [Pavan et al. \(2014\)](#). □

With (4) and the usual integration by parts, we can rewrite objective (3) as

$$\max_{(x,t) \in \mathcal{X}} \mathbb{E}[\theta_1 x_1(\theta_1) + \delta \theta_2 x_2(\theta_1, \theta_2)] - \mathbb{E} \left[\frac{1 - F_1(\theta_1)}{f_1(\theta_1)} (x_1(\theta_1) + \delta I(\theta_1, \theta_2) x_2(\theta_1, \theta_2)) \right] - V_{x,t}(\underline{\theta}_1).$$

Importantly, with a price mechanism, the seller cannot always extract all the expected surplus of the lowest type $\underline{\theta}_1$. In particular, the lowest type $\underline{\theta}_1$ can simply reject the first-period offer and obtain a price p_R for the second-period item, which implies that

$$V_{x,t}(\underline{\theta}_1) \geq \delta \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1].$$

With these observations, we reach the final form of our relaxation

$$\max_{(x,t) \in \mathcal{X}} \mathbb{E}[\varphi(\theta_1) x_1(\theta_1)] + \delta \mathbb{E}[\psi(\theta_1, \theta_2) x_2(\theta_1, \theta_2)] - \delta \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1] \quad (5)$$

where

$$\varphi(\theta_1) := \theta_1 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)}, \quad \psi(\theta_1, \theta_2) := \theta_2 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I(\theta_1, \theta_2)$$

are the virtual value functions for the first and second period, respectively. Note that if a mechanism (x^*, t^*) solves (5) and satisfies $V_{x^*, t^*}(\underline{\theta}_1) = \delta \mathbb{E}[(\theta_2 - p_R^*)_+ | \theta_1 = \underline{\theta}_1]$, then the mechanism also solves (3) and the value of (5) is exactly its expected revenue.

Step 3: To solve (5), note that under the constraints, the relevant parameters for the allocation rule are just (k, p_A, p_R) with $p_A \geq p_R$. We optimize directly over these:

$$\max_{k, p_A, p_R: p_A \geq p_R} \mathbb{E}[\varphi(\theta_1) \mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2) (\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R})] - \delta \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1].$$

Fix any k and consider the subproblem:

$$\max_{p_A, p_R: p_A \geq p_R} \mathbb{E}[\psi(\theta_1, \theta_2) (\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R})] - \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1]. \quad (6)$$

Define $D(\theta_1) := \inf_{\theta_2} \{\psi(\theta_1, \theta_2) \geq 0\}$.¹² If we allow for any allocation rule and the objective only consists of the first term, then we would set $x_2(\theta_1, \theta_2) = 1$ whenever $\psi(\theta_1, \theta_2) \geq 0$. That is, the second item is only allocated to types that fall above the boundary curve D in the type space $[\underline{\theta}_1, \bar{\theta}_1] \times [\underline{\theta}_2, \bar{\theta}_2]$. However, this is not a feasible solution to the constrained problem (6). To proceed, we make the following key observation:

¹²Moreover, let $D(\theta_1) := \bar{\theta}_2$ if $\psi(\theta_1, \theta_2) < 0$ for all θ_2 on which $\psi(\theta_1, \cdot)$ is defined; let $D(\theta_1) := \underline{\theta}_2$ if $\psi(\theta_1, \theta_2) \geq 0$ for all θ_2 on which $\psi(\theta_1, \cdot)$ is defined.

Claim 1. For any k , there exists an optimal solution to (6) with $p_A = p_R$.

Proof of Claim 1. Let $z = D(k)$ be the point where D crosses vertical line at k . Note that [Assumption A3](#) implies that $D(\cdot)$ is a non-increasing function. The existence of an optimal solution follows from standard compactness arguments. Fix any optimal solution (p_A, p_R) to program (6). If $p_A = p_R$, then we are done. Otherwise, there are three cases.

Case (i): $p_R < z < p_A$. As [Figure 2](#) shows, by lowering p_A the seller captures only types with $\psi(\theta_1, \theta_2) \geq 0$. Similarly, by raising p_R , the seller eliminates types with $\psi(\theta_1, \theta_2) < 0$. Importantly, the second term in the objective, $-\mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1]$, is also non-decreasing in p_R . Therefore, (z, z) must be weakly better than (p_A, p_R) , and hence (z, z) is an optimal solution.

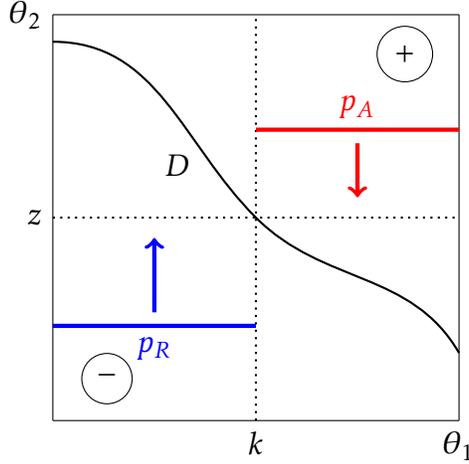


Figure 2: Allocation rule in the type space

Case (ii): $z \leq p_R < p_A$. By the same reasoning, lowering p_A to p_R weakly increases the objective. Case (iii): $p_R < p_A \leq z$. By the same reasoning, raising p_R to p_A weakly increases the objective. Thus, there exists an optimal solution with $p_A = p_R$. \square

Because [Claim 1](#) holds for any k , we can further reduce (5) to

$$\max_{k, p_2} \mathbb{E}[\varphi(\theta_1) \mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2) \mathbb{1}_{\theta_2 \geq p_2}] - \delta \mathbb{E}[(\theta_2 - p_2)_+ | \theta_1 = \underline{\theta}_1]. \quad (7)$$

Let (k^*, p_2^*) be an optimal solution to (7), which exists by compactness arguments. Consider the mechanism that posts prices k^*, p_2^* for the first-period and second-period items respectively. It implements the allocation rule $x_1(\theta_1) = \mathbb{1}_{\theta_1 \geq k^*}, x_2(\theta_1, \theta_2) = \mathbb{1}_{\theta_2 \geq p_2^*}$ that solves (7) and hence solves (5). In addition, the lowest type gets an expected payoff $V(\underline{\theta}_1) = \delta \mathbb{E}[(\theta_2 - p_2^*)_+ | \theta_1 = \underline{\theta}_1]$. Therefore, it solves (3) and its expected revenue is an

upper bound on what the seller can achieve in any pure-strategy PBE-D. Of course, its expected revenue is no higher than $\sum_t \delta^{t-1} \pi_t^*$, which completes the proof. *Q.E.D.*

Remark 3. In [Appendix A.3](#), we extend this proof to cover all PBE-D (including the ones in mixed strategies). In Step 1, we show that the first-period allocation is still a threshold rule and that the distribution of p_A must dominate the distribution of p_R in the sense that any realized p_A is weakly higher than any realized p_R . In Step 2, we consider a similar constrained mechanism design problem with the difference that the seller now chooses two distributions subject to a disjoint support constraint. In Step 3, we use [Claim 1](#) to show that the optimal value of the relaxed problem is again attained by simply posting two deterministic prices, one for each period.

Remark 4. Note that, as the proof shows, our result also holds in the setting where the seller can commit to any dynamic price mechanism (p_1, p_A, p_R) subject to the constraint $p_A \geq p_R$. Therefore, as indicated earlier, the only way that dynamic pricing can be profitable in our setting is to introduce repeated purchase discounts by setting $p_A < p_R$. Whenever offering repeated purchase discounts is not a feasible strategy (as in the examples discussed in the introduction), our proof shows that the seller cannot do better than simply posting all prices in advance.

Remark 5. Moreover, as the proof shows, our result also holds under any other equilibrium refinement that leads to the seller's belief of θ_2 dominating the prior in the strict likelihood ratio order after observing an off-path acceptance, and being dominated by the prior in the strict likelihood ratio order after observing an off-path rejection.

Remark 6. When the buyer's values are negatively correlated across periods, it is possible to sustain $p_A < p_R$ in a PBE-D. Then, consistent with our intuition, the seller may do strictly better in equilibrium than posting all prices in advance. Indeed, the following example shows that with negatively correlated values, the ratchet effect need not appear, and dynamic pricing can be profitable even under limited commitment.

Example 3 (Negative correlation). Suppose $T = 2$, $\Theta = \{1, 2\}$ and $\mathbb{P}(\theta_t = 1) = \mathbb{P}(\theta_t = 2) = 0.5$ for both periods. Suppose the types are perfectly anti-correlated. An optimal monopoly price is 2, and posting this price for both periods generates revenue 2. Consider instead the following strategy profile. The seller posts $p_1 = 2$ and then posts $p_A = 1$ following an acceptance, $p_R = 2$ following a rejection. Type (2, 1) accepts in both periods. Type (1, 2) rejects in period 1 and accepts in period 2. There is no profitable deviation for the buyer. This is also sequentially rational for the seller because she knows $\theta_2 = 2$ after a rejection and $\theta_2 = 1$ after an acceptance. Therefore, a seller-optimal PBE-D gives the seller revenue at least $2.5 > 2$.

4 Extensions

4.1 Complements

Our benchmark model assumes that the buyer's utility in the second period does not depend on the consumption of the first-period product, which rules out the possibility that the products may be complements. This is especially relevant in the setting of dynamic pricing, as one may expect that purchase histories become even more important.

However, we show that our main result continues to hold even if the products are complements. This is because the ratchet effect becomes stronger precisely when the past information becomes more relevant. When the products are substitutes, we provide an example to show that the ratchet effect need not appear, and the seller may do strictly better in equilibrium than posting all prices in advance (see [Remark 8](#) and [Example 4](#)).

To allow for complementarity, we let the distribution of θ_2 depend on the allocation of the first-period product. Specifically, the game proceeds exactly as in [Section 3](#), but $\theta_2 \sim F_2(\cdot|\theta_1, x_1)$ given the realization of θ_1 and the first-period allocation $x_1 \in \{0, 1\}$ (see [Remark 7](#) for an alternative formulation).

To capture complementarity, we assume that the distribution shifts upward in the sense of likelihood-ratio dominance after consuming the first product.

Assumption A4 (Complement). For any θ_1 , the conditional likelihood ratio $\frac{f(\theta_2|\theta_1, x_1=1)}{f(\theta_2|\theta_1, x_1=0)}$ is non-decreasing in θ_2 .

We adapt [Assumptions A1](#) to [A3](#) in [Section 3](#) to this setting.

Assumption A1' (MLRP). For any $\theta'_1 > \theta_1$ and $x_1 \in \{0, 1\}$, the conditional likelihood ratio $\frac{f(\theta_2|\theta'_1, x_1)}{f(\theta_2|\theta_1, x_1)}$ is strictly increasing in θ_2 .

Assumption A2' ($1/\delta$ -Lipschitz). For any $\theta'_1 > \theta_1$,

$$\mathbb{E}[\theta_2|\theta'_1, x_1 = 0] - \mathbb{E}[\theta_2|\theta_1, x_1 = 0] < \frac{1}{\delta}(\theta'_1 - \theta_1).$$

Assumption A3' (Regularity). The second-period virtual value function

$$\psi(\theta_1, \theta_2, x_1) := \theta_2 - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} I(\theta_1, \theta_2, x_1)$$

is non-decreasing in θ_1 , θ_2 , and x_1 , where $I(\theta_1, \theta_2, x_1) = -\frac{\partial F_2(\theta_2|\theta_1, x_1)}{\partial \theta_1} / f_2(\theta_2|\theta_1, x_1)$.

If the distribution of θ_2 conditional on θ_1 does not depend on x_1 as in [Section 3](#), then [Assumption A4](#) trivially holds and [Assumptions A1'](#) to [A3'](#) reduce to [Assumptions A1](#)

to [A3](#). In [Assumption A3'](#), the virtual value function being non-decreasing in x_1 is equivalent to $I(\theta_1, \theta_2, 1) \leq I(\theta_1, \theta_2, 0)$. Intuitively, this condition says that a perturbation in θ_1 would lead to a smaller change in θ_2 if the first product is consumed.

For example, suppose the type process follows an AR(1) process given first-period allocation $x \in \{0, 1\}$: $\theta_2 = \alpha\theta_1 + \epsilon_x$, where ϵ_x is drawn from a Gaussian distribution $\mathcal{N}(\mu_x, \sigma)$. Then [Assumption A4](#) is satisfied if $\mu_1 \geq \mu_0$. [Assumptions A1'](#) and [A2'](#) amount to $0 < \alpha < 1/\delta$. [Assumption A3'](#) is equivalent to that $F_1(\theta_1)$ has monotone hazard rate.

Theorem 2. Under [Assumptions A1'](#) to [A3'](#) and [A4](#),

- (i) there exists a PBE-D;
- (ii) in every PBE-D, the seller's revenue is no higher than what she would get by posting prices p_1, p_2 optimally in advance.

Proof of Theorem 2. See [Appendix A.6](#) for the proof of part (i). For expositional convenience, we prove part (ii) for all pure-strategy PBE-D here, and extend the proof to cover all PBE-D in [Appendix A.4](#).

As in [Section 3.1](#), we first show that in any equilibrium, the buyer follows a threshold rule. For that, it suffices to show that

$$h(\theta_1) := \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1, x_1 = 0] - \mathbb{E}[(\theta_2 - p_A)_+ | \theta_1, x_1 = 1]$$

satisfies $h(\theta'_1) - h(\theta_1) < \frac{1}{\delta}(\theta'_1 - \theta_1)$ for any $\theta'_1 > \theta_1$, which is proved in [Lemma 3](#) in [Appendix A.4](#). As shown in [Lemma 4](#) in [Appendix A.4](#), given that the preferences are persistent and the products are complements, after observing acceptance of the first-period offer, the seller believes that the buyer has a higher θ_2 and posts a higher price in the second period (i.e., $p_A \geq p_R$).

Then, as in [Section 3.1](#), using Theorem 1 in [Pavan et al. 2014](#), we can provide an upper bound for the seller's revenue by a constrained mechanism design problem with the price and allocation monotonicity constraints:

$$\begin{aligned} \max_{k, p_A, p_R: p_A \geq p_R} & \mathbb{E}[\varphi(\theta_1) \mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2, 0) \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R} | x_1 = 0] \\ & + \delta \mathbb{E}[\psi(\theta_1, \theta_2, 1) \mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} | x_1 = 1] - \delta \mathbb{E}[(\theta_2 - p_R)_+ | \underline{\theta}_1, x_1 = 0]. \end{aligned} \quad (8)$$

Note that both the virtual value function $\psi(\theta_1, \theta_2, x_1)$ and the distribution of θ_2 depend on the first-period allocation x_1 . Let D_0, D_1 be the boundary curves at which $\psi(\cdot, \cdot, 0), \psi(\cdot, \cdot, 1)$ cross zero respectively, as defined in [Section 3.1](#). Under the regularity condition [Assumption A3'](#), both $D_0(\theta_1)$ and $D_1(\theta_1)$ are non-increasing, and $D_1(\theta_1) \leq D_0(\theta_1)$ for all θ_1 .

For any cutoff k , define the effective boundary $D(\theta_1, k) := \mathbb{1}_{\theta_1 < k} D_0(\theta_1) + \mathbb{1}_{\theta_1 \geq k} D_1(\theta_1)$. Figure 3 illustrates. Note that $D(\theta_1, k)$ is non-increasing in θ_1 , for any given k . Therefore, by the proof of Claim 1, there exists an optimal solution to (8) with $p_A = p_R$.

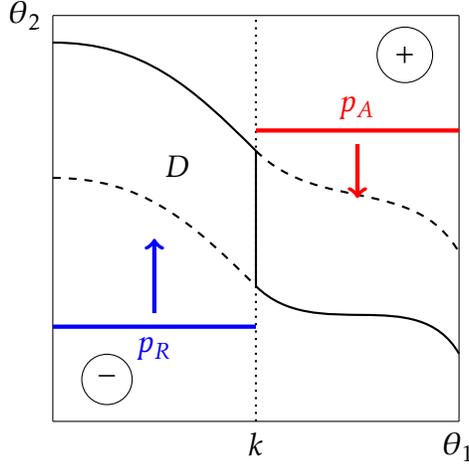


Figure 3: Allocation rule in the type space with the effective boundary

Let (k^*, p_2^*, p_2^*) be such an optimal solution to (8). Consider the mechanism that posts a price $p_1^* = k^* + \delta \mathbb{E}[(\theta_2 - p_2^*)_+ | k^*, x_1 = 1] - \delta \mathbb{E}[(\theta_2 - p_2^*)_+ | k^*, x_1 = 0]$ for the first-period item, and a price p_2^* for the second-period item. Note that type $\theta_1 = k^*$ is indifferent between accepting and rejecting the first-period offer. By Lemma 3, the first-period offer is accepted if and only if $\theta_1 \geq k^*$. Clearly, the second-period offer would be accepted if and only if $\theta_2 \geq p_2^*$. Thus, this mechanism implements the allocation rule $x_1(\theta_1) = \mathbb{1}_{\theta_1 \geq k^*}$, $x_2(\theta_1, \theta_2) = \mathbb{1}_{\theta_2 \geq p_2^*}$ that solves (8). Moreover, the lowest type gets an expected payoff $V(\underline{\theta}_1) = \delta \mathbb{E}[(\theta_2 - p_2^*)_+ | \underline{\theta}_1, x_1 = 0]$. Therefore, by the same argument as in Section 3.1, the seller's revenue in any PBE-D is no higher than what she would get from posting p_1, p_2 optimally in advance. \square

Remark 7. Instead of letting the distribution of θ_2 depend on x_1 , one may also introduce complementary products by assuming a non-additively separable utility for the buyer, $\theta_1 x_1 + \delta(\theta_2 x_2 + \kappa(\theta_1, \theta_2)x_1 x_2)$, where κ governs the extent of complementarity. Note that this is equivalent to our formulation. To see the equivalence, define $\tilde{\theta}_2 = \theta_2$ if $x_1 = 0$ and $\tilde{\theta}_2 = \theta_2 + \kappa(\theta_1, \theta_2)$ if $x_1 = 1$. Then, the buyer's utility is additively separable in $(\theta_1, \tilde{\theta}_2)$ and the distribution of $\tilde{\theta}_2$ depends on θ_1 and x_1 .

Remark 8. If the two items are substitutes, then consuming the first item decreases the demand for the second item. In that case, it is possible to sustain repeated purchase discounts (i.e., $p_A < p_R$) in equilibrium. Then, consistent with our intuition, the seller

may do strictly better in equilibrium than posting all prices in advance, as shown by the following example.

Example 4 (Substitutes). Suppose $\Theta = \{1, 2\}$ and $\mathbb{P}(\theta_t = 1) = \mathbb{P}(\theta_t = 2) = 0.5$ for both periods. Suppose that types are perfectly correlated and that the buyer's utility is given by $\theta_1 x_1 + \theta_2 x_2 - 1.5 x_1 x_2$ (see [Remark 7](#)). Since $\kappa(\theta_1, \theta_2) = -1.5$, the two goods are substitutes. If the seller posts p_1, p_2 in advance, then the seller can obtain at most revenue 1 (by posting $p_1 = p_2 = 1$). Consider instead the following strategy. The seller posts a price $p_1 = 1$ in the first period, then posts a price $p_A = 0.5$ after an acceptance and price $p_R = 1$ after a rejection. Both types accept in the first period, and the high type also accepts in the second period. Let the seller's belief be concentrated at $\theta_1 = 1$ after an off-path rejection. This strategy is then sequentially rational for the seller. Therefore, a seller-optimal PBE-D gives the seller revenue at least $1.25 > 1$.

4.2 Multiple Periods

Our benchmark model assumes that $T = 2$. One natural extension would call for a similar revenue bound for any finite time horizon T . With $T > 2$, however, it is known from [Devanur et al. \(2019\)](#) that a pure-strategy threshold equilibrium generally does not exist with perfectly correlated types (see Theorem 1 in [Devanur et al. 2019](#)).

To proceed, we introduce the following multiple-period game. In each period $t \in \{1, \dots, T\}$, the seller can either set the price p_t for period t only, or (publicly) commit to a sequence of prices $\{p_t, \dots, p_T\}$ for the unsold items. The buyer then decides whether to purchase the item as in [Section 3](#).

When $T = 2$, the modification reduces to providing the seller an option to post all prices in advance. So our main result is equivalent to that there exists a PBE-D of the modified game in which the seller simply posts all prices in advance.

We show that this statement (in fact, a stronger statement) extends to any finite time horizon T , under the following distributional assumption:

Assumption A5 (Log-concave AR(1)). The first-period type θ_1 follows a log-concave distribution F_1 (with bounded support). For all $t \geq 2$,

$$\theta_t = \alpha_t \theta_{t-1} + \epsilon_t$$

where $\alpha_t \in (0, \frac{1}{2\delta})$, and ϵ_t is independently drawn from a log-concave distribution G_t (with bounded support).

Recall that we say two equilibria are outcome-equivalent if they almost surely induce the same allocation and the same total discounted payment.

Theorem 3. Under [Assumption A5](#), for any finite T ,

- (i) there exists a PBE-D of the modified game;
- (ii) every PBE-D is outcome-equivalent to one in which the seller commits to $\{p_1^*, \dots, p_T^*\}$, where p_t^* is the monopoly price with respect to the marginal distribution F_t .

Proof of Theorem 3. See [Appendix A.7](#) for the proof of part (i). We now prove part (ii).

In [Lemma 5](#) in [Appendix A.5](#), we show that at any history \hat{h}_t , the buyer always adopts a threshold rule: $x_t(\hat{h}_t) = \mathbb{1}_{\theta_t \geq k_t}$, where the threshold k_t depends only on the public history. This statement is proved inductively by bounding the derivative of the buyer's continuation payoffs following an acceptance or a rejection and then showing that an appropriate single-crossing condition holds.

The rest of the proof proceeds by induction on T . The base case $T = 1$ is trivial. For the inductive step, fix any PBE-D. Consider what action the seller decides to take in period 2. Because the buyer adopts a threshold rule in period 1 (say the cutoff type is k), the seller's posterior belief about θ_1 follows a truncated distribution of F_1 after observing an acceptance or a rejection. Because log-concavity is preserved under truncation, convolution, and linear operations, her posterior belief about $(\theta_2, \theta_3, \dots, \theta_T)$ in period 2 also satisfies [Assumption A5](#).¹³ Moreover, because the buyer uses a threshold strategy, information about θ_1 does not matter to both players. Invoking the inductive hypothesis then implies that any equilibrium of the subgame must be outcome-equivalent to the one where the seller posts the monopoly prices from period 2 onward.

Let $\{p_t^A\}_{t=2}^T$ denote the sequence of monopoly prices with respect to the marginals $\{F_t(\cdot | \theta_1 \geq k)\}_{t=2}^T$ and similarly define $\{p_t^R\}_{t=2}^T$. Because the posterior distributions are all log-concave, both $\{p_t^A\}_{t=2}^T$ and $\{p_t^R\}_{t=2}^T$ are uniquely defined. Moreover, because these prices are unique, it is without loss of generality to consider only pure strategies. Note that under [Assumption A5](#), (θ_1, θ_t) satisfies an appropriate MLRP condition (see [Lemma 6](#) in [Appendix A.5](#)). Therefore, we have $p_t^A \geq p_t^R$ for all t (see [Lemma 7](#) in [Appendix A.5](#)). This enables us to adopt the same approach as in the proof of [Theorem 1](#). In particular, using [Theorem 1](#) of [Pavan et al. \(2014\)](#), we have that the following relaxed problem gives

¹³When $k = \underline{\theta}_1$ (or $\bar{\theta}_1$), as noted in [Remark 1](#) and shown in [Appendix B](#), it is without loss of generality to let the seller's belief of θ_1 be $\underline{\theta}_1$ (or $\bar{\theta}_1$) when observing an off-path rejection (or acceptance).

an upper bound for the seller's revenue:

$$\max_{(x,t) \in \mathcal{X}} \mathbb{E} \left[\sum_{t=1}^T \delta^{t-1} \psi_t(\theta_1, \dots, \theta_t) x_t(\theta_1, \dots, \theta_t) \right] - V_{x,t}(\underline{\theta}_1) \quad (9)$$

where \mathcal{X} encodes that the seller selects a price mechanism subject to the threshold allocation constraint for period 1 and price monotonicity constraint for period $t \geq 2$, and ψ_t is the period- t virtual value function (see Section 4 of [Pavan et al. 2014](#)). In particular, under [Assumption A5](#), we have

$$\psi_t(\theta_1, \dots, \theta_t) = \theta_t - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} \prod_{s=2}^t \alpha_s$$

(see Example 2 in [Pavan et al. 2014](#)). By the same argument as in [Section 3.1](#), the problem reduces to

$$\max_{\substack{k, \{p_t^A\}_{t=2}^T, \{p_t^R\}_{t=2}^T \\ p_t^A \geq p_t^R, \forall t \geq 2}} \mathbb{E}[\varphi(\theta_1) \mathbb{1}_{\theta_1 \geq k}] + \sum_{t=2}^T \delta^{t-1} \left(\mathbb{E}[\psi_t(\theta_1, \theta_t) (\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_t \geq p_t^A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_t \geq p_t^R})] - \mathbb{E}[(\theta_t - p_t^R)_+ | \underline{\theta}_1] \right).$$

For any k , the t -th term in the above sum is of the same form as in (6). For any $t \geq 2$, note that $\psi_t(\theta_1, \theta_t)$ is strictly increasing in θ_t and non-decreasing in θ_1 because log-concavity implies monotone hazard rate. Thus, by [Lemma 8](#) in [Appendix A.5](#), if (p_t^A, p_t^R) induces a period- t allocation rule that is not almost surely equal to $\mathbb{1}_{\theta_t \geq p_t}$ for some p_t , then the seller does strictly worse for the period- t revenue. Then, for any k , the outcome under any optimal choice of $(\{p_t^A\}_{t=2}^T, \{p_t^R\}_{t=2}^T)$ must be equivalent to the outcome under some price path $\{p_2, \dots, p_T\}$. Therefore, for the T -period game, every PBE-D must be outcome-equivalent to one in which the seller commits to $\{p_1^*, \dots, p_T^*\}$. This proves the inductive step. \square

5 Conclusion

This paper studies dynamic pricing with persistent private information and limited commitment. In a two-period game, the main result shows that regardless of the degree of persistence of the private information, any equilibrium under a D1-style refinement gives the seller revenue no higher than what she would get from posting all prices in advance. The ability to condition prices on purchase histories, which are informative about the buyer's values, introduces the ratchet effect that is strong enough to outweigh the gains

from learning.

This result contrasts with the case of full commitment — a large literature on dynamic mechanism design highlights the power of dynamically responding to the buyer’s evolving private information. As our analysis shows, to benefit from dynamic pricing, the seller must commit to a lower price precisely when he *ex post* knows that the buyer has a higher willingness to pay. That is, dynamic rent extraction always goes in the opposite direction of the seller’s sequential rationality, and thus hinges critically on her long-term commitment power. Building on this insight, we also show how similar results hold when the products are complements or when there are multiple periods. In contrast, when the products are substitutes or when the values are negatively correlated, the ratchet effect need not appear. Then, as we show by example, dynamic pricing can be effective even under limited commitment.

Appendix A Omitted Proofs

A.1 Proof of Lemma 1

Fix any $k \in (\underline{\theta}_1, \bar{\theta}_1)$. Note that

$$\frac{f(\theta_2|\theta_1 \geq k)}{f(\theta_2|\theta_1 < k)} \propto \frac{\int_k^{\bar{\theta}_1} f(\theta_2|\theta_1 = s')f(\theta_1 = s')ds'}{\int_{\underline{\theta}_1}^k f(\theta_2|\theta_1 = s)f(\theta_1 = s)ds}.$$

It then suffices to show that for any $s' > k$,

$$\frac{f(\theta_2|\theta_1 = s')}{\int_{\underline{\theta}_1}^k f(\theta_2|\theta_1 = s)f(\theta_1 = s)ds}$$

is strictly increasing in θ_2 . Equivalently,

$$\frac{\int_{\underline{\theta}_1}^k f(\theta_2|\theta_1 = s)f(\theta_1 = s)ds}{f(\theta_2|\theta_1 = s')}$$

is strictly decreasing in θ_2 . It then suffices to show

$$\frac{f(\theta_2|\theta_1 = s)}{f(\theta_2|\theta_1 = s')}$$

is strictly decreasing in θ_2 for every $s' > s$, but that is implied by the MLRP property.

A.2 Proof of Lemma 2

It is clear that

$$p_i \in \operatorname{argmax}_p [\log(p) + \log(1 - F_i(p))].$$

The claim follows from the monotone selection theorem, if we show that $v(p, i) := \log(1 - F_i(p))$ has strictly increasing differences. Let

$$\Delta(p) := \log(1 - F_2(p)) - \log(1 - F_1(p)).$$

Note that

$$\Delta'(p) = \frac{f_1(p)}{1 - F_1(p)} - \frac{f_2(p)}{1 - F_2(p)} > 0$$

since strict likelihood-ratio dominance implies strict hazard-rate dominance.

A.3 Completion of Proof of Part (ii) of Theorem 1

We complete the proof of part (ii) of Theorem 1 by extending the proof in the main text to cover all PBE-D. Instead of choosing two fixed prices p_A, p_R , the seller now chooses two random variables \tilde{p}_A, \tilde{p}_R . We refer to their realizations as p_A, p_R . (It is without loss of generality to only focus on mixing over p_A, p_R since mixing over p_1 implies that there always exists a deterministic p_1 achieving the same revenue.)

Step 1: We show that the first-period allocation rule is still a threshold rule, and that price monotonicity holds in the sense that $p_A \geq p_R$ for any realized p_A, p_R , i.e., $\min\{\operatorname{supp}(\tilde{p}_A)\} \geq \max\{\operatorname{supp}(\tilde{p}_R)\}$. Note that (2) is preserved under averaging. So if we define

$$h(\theta_1) := \mathbb{E}_{\tilde{p}_A, \tilde{p}_R} [\mathbb{E}[(\theta_2 - p_R)_+ - (\theta_2 - p_A)_+ | \theta_1]]$$

then we still have $h(\theta'_1) - h(\theta_1) < \frac{1}{\delta}(\theta'_1 - \theta_1)$ for any $\theta'_1 > \theta_1$. Type θ_1 accepts if

$$(\theta_1 - p_1) + \delta \mathbb{E}_{\tilde{p}_A} [\mathbb{E}[(\theta_2 - p_A)_+ | \theta_1]] \geq \delta \mathbb{E}_{\tilde{p}_R} [\mathbb{E}[(\theta_2 - p_R)_+ | \theta_1]].$$

Then as before, $g(\theta_1) := \frac{1}{\delta}(\theta_1 - p_1)$ crosses h at most once from below. Therefore, the first-period allocation is characterized by a cutoff type k .

Suppose $k \in (\underline{\theta}_1, \bar{\theta}_1)$. By Lemma 1, $F_2(\cdot | \theta_1 \geq k) \succ_L F_2(\cdot | \theta_1 < k)$. Note that any realized p_A must be an optimal price under belief $F_2(\cdot | \theta_1 \geq k)$, as the seller only mixes among optimal prices. Similarly, any realized p_R must be an optimal price under belief $F_2(\cdot | \theta_1 <$

k). Hence $p_A \geq p_R$ for any realized p_A, p_R by [Lemma 2](#). The case of $k = \underline{\theta}_1$ or $\bar{\theta}_1$ is handled similarly as in the proof of [Proposition 1](#).

Step 2: The same relaxation argument applies. The only difference is that the seller now commits to a randomized price mechanism in the second period with the constraint

$$\min\{\text{supp}(\tilde{p}_A)\} \geq \max\{\text{supp}(\tilde{p}_R)\}$$

denoted by $\tilde{p}_A \geq \tilde{p}_R$. The same steps go through, except that the relaxed problem is now

$$\max_{k, \tilde{p}_A, \tilde{p}_R: \tilde{p}_A \geq \tilde{p}_R} \mathbb{E}[\varphi(\theta_1)\mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2)(\mathbb{1}_{\theta_1 \geq k}\mathbb{1}_{\theta_2 \geq \tilde{p}_A} + \mathbb{1}_{\theta_1 < k}\mathbb{1}_{\theta_2 \geq \tilde{p}_R})] - \delta \mathbb{E}[(\theta_2 - \tilde{p}_R)_+ | \theta_1 = \underline{\theta}_1]$$

where \tilde{p}_A, \tilde{p}_R are two random variables.

Step 3: To solve this, we consider a further relaxation. We couple \tilde{p}_A, \tilde{p}_R in a common probability space (Ω, \mathcal{F}, Q) . Consider the requirement that $\tilde{p}_A(\omega) \geq \tilde{p}_R(\omega)$ for all $\omega \in \Omega$, denoted by $\tilde{p}_A \geq \tilde{p}_R$. The disjoint support constraint clearly implies this, so this is a relaxation of the problem above. The seller's problem is to maximize over k and two \mathcal{F} -measurable functions:

$$\max_{k, \tilde{p}_A, \tilde{p}_R: \tilde{p}_A \geq \tilde{p}_R} \mathbb{E}_Q \left[\mathbb{E}[\varphi(\theta_1)\mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2)(\mathbb{1}_{\theta_1 \geq k}\mathbb{1}_{\theta_2 \geq \tilde{p}_A(\omega)} + \mathbb{1}_{\theta_1 < k}\mathbb{1}_{\theta_2 \geq \tilde{p}_R(\omega)})] - \delta \mathbb{E}[(\theta_2 - \tilde{p}_R(\omega))_+ | \theta_1 = \underline{\theta}_1] \right].$$

Now fix any ω . Let p_A, p_R denote $\tilde{p}_A(\omega), \tilde{p}_R(\omega)$. Then the ex post maximization is

$$\max_{k, p_A, p_R: p_A \geq p_R} \mathbb{E}[\varphi(\theta_1)\mathbb{1}_{\theta_1 \geq k}] + \delta \mathbb{E}[\psi(\theta_1, \theta_2)(\mathbb{1}_{\theta_1 \geq k}\mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k}\mathbb{1}_{\theta_2 \geq p_R})] - \delta \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1 = \underline{\theta}_1]$$

which is exactly what we have in [Section 3.1](#) and has a maximizer (k^*, p_2^*, p_2^*) , irrespective of ω . Hence, there is a solution to the ex ante maximization problem (with the constraint $\tilde{p}_A \geq \tilde{p}_R$) with $\tilde{p}_A = \tilde{p}_R = p^*$. This solution also satisfies the constraint $\tilde{p}_A \geq \tilde{p}_R$ and thus solves the mechanism design relaxation where the seller is allowed to use randomized price mechanisms in the second period. So the seller's optimal constrained mechanism is still to simply post two deterministic prices in advance. Therefore, her revenue in any PBE-D is no higher than $\sum_t \delta^{t-1} \pi_t^*$.

A.4 Completion of Proof of Part (ii) of Theorem 2

Lemma 3. *In every pure-strategy PBE-D, $h(\theta'_1) - h(\theta_1) < \frac{1}{\delta}(\theta'_1 - \theta_1)$ for any $\theta'_1 > \theta_1$.*

Proof. Fix any $\theta'_1 > \theta_1$. Integration by parts implies

$$\begin{aligned} & \mathbb{E}[(\theta_2 - p_R)_+ | \theta'_1, x_1 = 0] - \mathbb{E}[(\theta_2 - p_R)_+ | \theta_1, x_1 = 0] \\ &= \int_0^\infty [\mathbb{P}(\theta_2 > s | \theta'_1, x_1 = 0) - \mathbb{P}(\theta_2 > s | \theta_1, x_1 = 0)] \mathbb{1}_{s \geq p_R} ds \\ &\leq \int_0^\infty [\mathbb{P}(\theta_2 > s | \theta'_1, x_1 = 0) - \mathbb{P}(\theta_2 > s | \theta_1, x_1 = 0)] ds \\ &= \mathbb{E}[\theta_2 | \theta'_1, x_1 = 0] - \mathbb{E}[\theta_2 | \theta_1, x_1 = 0] < \frac{1}{\delta}(\theta'_1 - \theta_1) \end{aligned}$$

where we used FOSD (implied by [Assumption A1'](#)) and [Assumption A2'](#). The claim follows by noting that FOSD implies $\mathbb{E}[(\theta_2 - p_A)_+ | \theta'_1, x_1 = 1] - \mathbb{E}[(\theta_2 - p_A)_+ | \theta_1, x_1 = 1] \geq 0$ \square

Lemma 4. *In every pure-strategy PBE-D, $p_A \geq p_R$.*

Proof. Using the assumptions, we have

$$F_2(\cdot | \theta'_1, x_1 = 1) \succ_L F_2(\cdot | \theta_1, x_1 = 1) \succeq_L F_2(\cdot | \theta_1, x_1 = 0)$$

for any $\theta'_1 > \theta_1$, where \succeq_L denotes (weak) likelihood-ratio dominance. This then implies that for any $k \in (\underline{\theta}_1, \bar{\theta}_1)$, $F_2(\cdot | \theta_1 \geq k, x_1 = 1) \succ_L F_2(\cdot | \theta_1 < k, x_1 = 0)$ by the proof of [Lemma 1](#). Price monotonicity then follows from [Lemma 2](#). When $k = \underline{\theta}_1$ or $\bar{\theta}_1$, the same holds by the D1 criterion as in the proof of [Proposition 1](#). \square

Extension to all PBE-D. We follow the same notation and the same steps as in [Appendix A.3](#). For Step 1, note that [Lemma 3](#) continues to hold if we define

$$h(\theta_1) := \mathbb{E}_{\bar{p}_A, \bar{p}_R} \left[\mathbb{E}[(\theta_2 - p_R)_+ | \theta_1, x_1 = 0] - \mathbb{E}[(\theta_2 - p_A)_+ | \theta_1, x_1 = 1] \right].$$

Moreover, by the same reasoning as in [Appendix A.3](#) and the proof of [Lemma 4](#), we have $p_A \geq p_R$ for any realized p_A, p_R . Steps 2 and 3 proceed in exactly the same way as in [Appendix A.3](#), and are thus omitted.

A.5 Completion of Proof of Part (ii) of Theorem 3

Lemma 5. *For any seller's strategy, the buyer's optimal strategy is given by a threshold rule: at every history \hat{h}_t , $x_t(\hat{h}_t) = \mathbb{1}_{\theta_t \geq k_t}$, where k_t depends only on the public history.*

Proof. Note that at any history \hat{h}_t , the buyer solves the following problem

$$\max_{x \in \{0,1\}} [(\theta_t - p_t) + U_t^A(\theta_t, h_t^\dagger)]x + U_t^R(\theta_t, h_t^\dagger)(1 - x)$$

where U_t^A, U_t^R are the continuation payoffs following an acceptance and a rejection, and we decompose the history of the buyer \hat{h}_t into (θ_t, h_t^\dagger) .

We prove the following three claims together by backward induction:

- (1) $x_t(\hat{h}_t) = \mathbb{1}_{\theta_t \geq k_t}$, where k_t depends only on the public history.
- (2) $U_t^A(\cdot, h_t^\dagger), U_t^R(\cdot, h_t^\dagger)$ are 1-Lipschitz and non-decreasing functions of θ_t .
- (3) $U_t^A(\theta_t, \cdot), U_t^R(\theta_t, \cdot)$ depend on h_t^\dagger only through public history $h_t^o := (h_t, p_t)$.

Base case: In the last period, $k_T = p_T$ and $U_T^A = U_T^R = 0$. So the claims hold.

Inductive step: Note that

$$U_t^A(\theta_t, h_t^\dagger) = \delta \mathbb{E} \left[\max \left\{ \theta_{t+1} - p_{t+1} + U_{t+1}^A(\theta_{t+1}, h_{t+1}^\dagger), U_{t+1}^R(\theta_{t+1}, h_{t+1}^\dagger) \right\} \middle| \theta_t, h_t^\dagger, x_t = A \right].$$

By definition, p_{t+1} depends only on the public history. Since the type process is Markovian, the distribution of θ_{t+1} depends only on θ_t . By the inductive hypothesis, $U_{t+1}^A(\theta_{t+1}, \cdot), U_{t+1}^R(\theta_{t+1}, \cdot)$ depend only on the public history. Therefore, for any fixed θ_t , $U_t^A(\theta_t, \cdot)$ depends on h_t^\dagger only through the public history h_t^o . The same argument works for $U_t^R(\theta_t, \cdot)$. Moreover, by the inductive hypothesis, for any fixed p ,

$$\max \left\{ \theta_{t+1} - p + U_{t+1}^A(\theta_{t+1}, h_{t+1}^o), U_{t+1}^R(\theta_{t+1}, h_{t+1}^o) \right\}$$

is 2-Lipschitz and non-decreasing in θ_{t+1} . Therefore,

$$U_{t+1}(\theta_{t+1}, h_t^o) := \mathbb{E}_{p_{t+1} \sim \sigma} \left[\max \left\{ \theta_{t+1} - p_{t+1} + U_{t+1}^A(\theta_{t+1}, h_{t+1}^o), U_{t+1}^R(\theta_{t+1}, h_{t+1}^o) \right\} \middle| h_t^o, x_t = A \right]$$

is also 2-Lipschitz and non-decreasing in θ_{t+1} , where σ is the seller's strategy. Fix any $\theta'_t > \theta_t$, integration by parts yields that

$$U_t^A(\theta'_t, h_t^o) - U_t^A(\theta_t, h_t^o) = \delta \int_0^\infty [\mathbb{P}(\theta_{t+1} > s | \theta'_t) - \mathbb{P}(\theta_{t+1} > s | \theta_t)] U'_{t+1}(s, h_t^o) ds.$$

Under [Assumption A5](#), $\mathbb{P}(\theta_{t+1} > s|\theta'_t) \geq \mathbb{P}(\theta_{t+1} > s|\theta_t)$. Thus,

$$\begin{aligned} U_t^A(\theta'_t, h_t^o) - U_t^A(\theta_t, h_t^o) &\leq 2\delta \int_0^\infty [\mathbb{P}(\theta_{t+1} > s|\theta'_t) - \mathbb{P}(\theta_{t+1} > s|\theta_t)] ds \\ &= 2\delta(\mathbb{E}[\theta_{t+1}|\theta'_t] - \mathbb{E}[\theta_{t+1}|\theta_t]) < \theta'_t - \theta_t \end{aligned}$$

where the last inequality is due to the assumption $\alpha_t \in (0, \frac{1}{2\delta})$ for all t . This shows that $U_t^A(\cdot, h_t^o)$ is 1-Lipschitz. It is clear from the above that $U_t^A(\cdot, h_t^o)$ is non-decreasing. The same arguments work for $U_t^R(\cdot, h_t^o)$. Together these also imply that for any h_t^\dagger ,

$$H(\theta_t, h_t^\dagger) := (\theta_t - p_t) + U_t^A(\theta_t, h_t^\dagger) - U_t^R(\theta_t, h_t^\dagger) = (\theta_t - p_t) + U_t^A(\theta_t, h_t^o) - U_t^R(\theta_t, h_t^o)$$

is a strictly increasing function and thus crosses 0 at most once from below. Therefore, the buyer's strategy is given by a threshold rule and the threshold k_t depends only on the public history h_t^o . This proves the inductive step. \square

Lemma 6. For any $2 \leq t \leq T$ and any $\theta'_1 > \theta_1$, the conditional likelihood ratio $\frac{f(\theta_t|\theta'_1)}{f(\theta_t|\theta_1)}$ is non-decreasing in θ_t .

Proof. Fix any t . We can write

$$\theta_t = \alpha_t \theta_{t-1} + \epsilon_t = \alpha_t(\alpha_{t-1} \theta_{t-2} + \epsilon_{t-1}) + \epsilon_t = \dots = \left(\prod_{s=2}^t \alpha_s \right) \theta_1 + \sum_{s=2}^t \left(\prod_{n=s+1}^t \alpha_n \right) \epsilon_s$$

where we interpret $\prod_{n=s+1}^t \alpha_n = 1$ when $s = t$. Let $\alpha := \prod_{s=2}^t \alpha_s > 0$. Since ϵ_s are mutually independent and each follows a log-concave distribution, $\epsilon := \sum_{s=2}^t (\prod_{n=s+1}^t \alpha_n) \epsilon_s$ follows a log-concave distribution. Moreover, ϵ and θ_1 are independent. Therefore $\theta_t = \alpha \theta_1 + \epsilon$ where $\alpha > 0$ and ϵ follows some log-concave distribution G (with a density g). Since G is a log-concave distribution, the conditional distribution of θ_t has MLRP in $\alpha \theta_1$ by standard results, i.e.,

$$\frac{f_t(\theta_t|\theta'_1)}{f_t(\theta_t|\theta_1)} = \frac{g(\theta_t - \alpha \theta'_1)}{g(\theta_t - \alpha \theta_1)}$$

is non-decreasing in θ_t , for any $\alpha \theta'_1 > \alpha \theta_1$. Since $\alpha > 0$, the latter is simply $\theta'_1 > \theta_1$, proving the claim. \square

Lemma 7. For any $2 \leq t \leq T$, $p_t^A \geq p_t^R$.

Proof. Because $F_t(\cdot|\theta_1 < k)$ and $F_t(\cdot|\theta_1 \geq k)$ are both log-concave, both admit a unique optimal monopoly price.¹⁴ Then, we no longer need the monotone selection theorem

¹⁴To see this, note that if F is log-concave, then $\log(p) + \log(1 - F(p))$ is strictly concave and admits a

used in [Lemma 2](#). It suffices to show that $F_t(\cdot|\theta_1 \geq k)$ weakly dominates $F_t(\cdot|\theta_1 < k)$ in hazard rate order, which is implied by weak likelihood-ratio dominance $F_t(\cdot|\theta_1 \geq k) \geq_L F_t(\cdot|\theta_1 < k)$. By the proof of [Lemma 1](#), it suffices to show (θ_1, θ_t) satisfies the weak MLRP condition, which is established in [Lemma 6](#). The case of $k = \underline{\theta}_1$ or $\bar{\theta}_1$ follows by a similar argument as in the proof of [Proposition 1](#). \square

Lemma 8. *Suppose $\psi(\theta_1, \theta_2)$ is non-decreasing in θ_1 and strictly increasing in θ_2 . Then, for any $k \in [\underline{\theta}_1, \bar{\theta}_1]$, and any optimal solution (p_A, p_R) to the following problem*

$$\max_{p_A, p_R: p_A \geq p_R} \mathbb{E}[\psi(\theta_1, \theta_2)(\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R})] - \mathbb{E}[(\theta_2 - p_R)_+ | \underline{\theta}_1],$$

there exists an optimal solution $(\tilde{p}_A, \tilde{p}_R)$ to the same problem such that (i) $\tilde{p}_A = \tilde{p}_R$ and (ii)

$$\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R} \stackrel{a.s.}{=} \mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq \tilde{p}_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq \tilde{p}_R}.$$

Proof. Let $D(\theta_1)$ be defined as in [Section 3.1](#) and $z := D(k)$. Note that $D(\cdot)$ is non-increasing. Fix any optimal solution with $p_A > p_R$. Suppose $p_R < z < p_A$. Then, the set of types $[\underline{\theta}_1, k] \times [p_R, z]$ must have measure 0 because otherwise increasing p_R up to z strictly increases the objective. Similarly, the set of types $[k, \bar{\theta}_1] \times [z, p_A]$ must also have measure 0. Thus, the second-period allocation rule $\mathbb{1}_{\theta_1 \geq k} \mathbb{1}_{\theta_2 \geq p_A} + \mathbb{1}_{\theta_1 < k} \mathbb{1}_{\theta_2 \geq p_R} \stackrel{a.s.}{=} \mathbb{1}_{\theta_2 \geq z}$. The same argument holds for the other two cases $z \leq p_R < p_A$ and $p_R < p_A \leq z$ as well. \square

A.6 Proof of Part (i) of [Theorem 1](#) and Part (i) of [Theorem 2](#)

As explained in [Remark 1](#) and shown formally in [Appendix B](#), it suffices to show the existence of a PBE with a particular class of off-path beliefs.

Definition 1. A PBE- \star is a perfect Bayesian equilibrium with the following restrictions:

- After observing an off-path rejection in period t under belief μ_t on θ_t , the seller believes that the period- t type of the buyer is $\min\{\text{supp}(\mu_t)\}$.
- After observing an off-path acceptance in period t under belief μ_t on θ_t , the seller believes that the period- t type of the buyer is $\max\{\text{supp}(\mu_t)\}$.

In what follows, we show the existence of a PBE- \star for the game described in [Section 4.1](#), which then implies part (i) of [Theorem 1](#) and part (i) of [Theorem 2](#).

unique maximizer.

Consider the following auxiliary complete information one-shot game parameterized by p with 3 players: B, S_A, S_R . Player B chooses $k \in [\underline{\theta}_1, \bar{\theta}_1]$ to maximize

$$\mathbb{E}[\{(\theta_1 - p) + \delta(\theta_2 - p_A)_+\} \mathbb{1}_{\theta_1 \geq k} | x_1 = 1] + \mathbb{E}[\delta(\theta_2 - p_R)_+ \mathbb{1}_{\theta_1 \leq k} | x_1 = 0].$$

Player S_A chooses $p_A \in [\underline{\theta}_2, \bar{\theta}_2]$ to maximize

$$(1 - F_2(p_A | \theta_1 \geq k, x_1 = 1)) p_A.$$

Player S_R chooses $p_R \in [\underline{\theta}_2, \bar{\theta}_2]$ to maximize

$$(1 - F_2(p_R | \theta_1 \leq k, x_1 = 0)) p_R.$$

Note that this is a continuous game. By [Glicksberg \(1952\)](#), there exists a mixed-strategy Nash equilibrium, denoted by measures $(\eta_B^p, \eta_A^p, \eta_R^p)$.

By [Lemma 3](#), we know that Player B will not mix. Thus we can identify the equilibrium strategy η_B^p by a cutoff k^p . For any fixed p , let \mathcal{E}^p be the set of equilibrium tuples (k^p, π_A^p, π_R^p) , where π_A^p, π_R^p are equilibrium payoffs of players S_A and S_R . We claim that \mathcal{E}^p is a closed and bounded set in \mathbb{R}^3 . Clearly, it is bounded. To show it is closed, take any sequence $(k_n^p, \pi_{A,n}^p, \pi_{R,n}^p) \in \mathcal{E}^p$ converging to (k^p, π_A^p, π_R^p) . Let $(\eta_{B,n}^p, \eta_{A,n}^p, \eta_{R,n}^p)$ be the corresponding equilibrium points. Note that this family is tight (each equilibrium point is viewed as a product measure on a compact subset of \mathbb{R}^3). Therefore, by Prokhorov's Theorem, this family converges weakly along a subsequence to some $(\eta_B^p, \eta_A^p, \eta_R^p)$. By Theorem 2 in [Milgrom and Weber \(1985\)](#), $(\eta_B^p, \eta_A^p, \eta_R^p)$ is an equilibrium point. Since the utility functions are bounded and continuous, the corresponding tuple for this equilibrium is (k^p, π_A^p, π_R^p) and therefore $(k^p, \pi_A^p, \pi_R^p) \in \mathcal{E}^p$. Now define the function

$$G(k, \pi_A, \pi_R; p) := (p + \delta\pi_A)(1 - F_1(k)) + \delta\pi_R F_1(k)$$

which is jointly continuous in (k, π_A, π_R) and hence attains its maximum on the compact set \mathcal{E}^p . Let

$$H(p) := \max_{(k, \pi_A, \pi_R) \in \mathcal{E}^p} G(k, \pi_A, \pi_R; p).$$

We claim that $H(p)$ is upper semi-continuous. To see this, fix any p_0 . Let $\{p_n\}_n$ converging to p_0 be a sequence along which we obtain $\limsup_{p \rightarrow p_0} H(p)$. Let

$$(k_n^*, \pi_{A,n}^*, \pi_{R,n}^*) \in \operatorname{argmax}_{(k, \pi_A, \pi_R) \in \mathcal{E}^{p_n}} G(k, \pi_A, \pi_R; p_n).$$

Let $(\eta_{B,n}^*, \eta_{A,n}^*, \eta_{R,n}^*)$ be the corresponding equilibrium points. By the same argument as before, along a subsequence we know that $(\eta_{B,n}^*, \eta_{A,n}^*, \eta_{R,n}^*)$ converges weakly to some $(\eta_B^*, \eta_A^*, \eta_R^*)$, which by [Milgrom and Weber \(1985\)](#) is an equilibrium point of the game with parameter p_0 . Then as before $(k_n^*, \pi_{A,n}^*, \pi_{R,n}^*)$ also converges (along the subsequence) to (k^*, π_A^*, π_R^*) given by the equilibrium $(\eta_B^*, \eta_A^*, \eta_R^*)$.¹⁵ By continuity of G ,

$$\limsup_{p \rightarrow p_0} H(p) = \lim_{n \rightarrow \infty} G(k_n^*, \pi_{A,n}^*, \pi_{R,n}^*; p_n) = G(k^*, \pi_A^*, \pi_R^*; p_0) \leq H(p_0)$$

where the last inequality is due to $(k^*, \pi_A^*, \pi_R^*) \in \mathcal{E}^{p_0}$, proving the claim.

Since H is upper semi-continuous, and the set $[-(\bar{\theta}_1 + \bar{\theta}_2), \bar{\theta}_1 + \bar{\theta}_2]$ is compact, there exists a p^* that maximizes $H(p)$. We now recover a PBE- \star of the original game. Let the seller offer p^* in period 1. For any p offered by the seller, let (k^p, π_A^p, π_R^p) be a tuple in \mathcal{E}^p that maximizes $G(\cdot; p)$. Let $(k^p, \eta_A^p, \eta_R^p)$ be the corresponding equilibrium point. Let the buyer use strategy $x_1(\theta_1, p) = \mathbb{1}_{\theta_1 \geq k^p}$. Following histories A, R , the seller offers prices according to the mixed strategies η_A^p, η_R^p , respectively. Let her beliefs be $\mu_A = F_2(\cdot | \theta_1 \geq k^p)$ and $\mu_R = F_2(\cdot | \theta_1 \leq k^p)$. The buyer accepts in the second period if and only if his type is above the price.

By inspection, this is a PBE of the original game. Note that the construction guarantees that whenever $k = \bar{\theta}_1$, η_A maximizes the seller's payoff with respect to $F_2(\cdot | \theta_1 = \bar{\theta}_1, x_1 = 1)$ and whenever $k = \underline{\theta}_1$, η_R maximizes the seller's payoff with respect to $F_2(\cdot | \theta_1 = \underline{\theta}_1, x_1 = 0)$. Thus, the PBE is in fact a PBE- \star .

A.7 Proof of Part (i) of [Theorem 3](#)

We show that there exists a PBE- \star of the subgame after the seller chooses not to commit in the first period, which immediately implies that the whole game also has a PBE- \star (see [Appendix A.6](#) and [Definition 1](#) for PBE- \star). The proof proceeds in two steps.

Step 1: We first show that the following game has a mixed-strategy PBE- \star : for a fixed $T \geq 2$, the seller posts a price p in period 1, and the buyer decides whether to accept; following histories A, R , the seller commits to prices $\{p_t^A\}_{t=2}^T, \{p_t^R\}_{t=2}^T$, respectively. The payoffs and the evolution of buyer's types are the same as in [Section 4.2](#).

Lemma 9. *There exists a PBE- \star of this game.*

¹⁵One can view $(\eta_{B,n}^*, \eta_{A,n}^*, \eta_{R,n}^*, p_n)$ as a product measure on some compact subset of \mathbb{R}^4 , and the claim follows by observing that the utility functions are bounded and jointly continuous in the actions and p .

Proof. Consider the following auxiliary complete information one-shot game with 3 players: B, S_A, S_R . Player B chooses $k \in [\underline{\theta}_1, \bar{\theta}_1]$ to maximize

$$\mathbb{E} \left[\left\{ (\theta_1 - p) + \sum_{t=2}^T \delta^{t-1} (\theta_t - p_t^A)_+ \right\} \mathbb{1}_{\theta_1 \geq k} \right] + \mathbb{E} \left[\sum_{t=2}^T \delta^{t-1} (\theta_t - p_t^R)_+ \mathbb{1}_{\theta_1 \leq k} \right].$$

Player S_A chooses $\{p_t^A\}_{t=2}^T \in \prod_{t=2}^T [\underline{\theta}_t, \bar{\theta}_t]$ to maximize

$$\sum_{t=2}^T \delta^{t-2} (1 - F_t(p_t^A | \theta_1 \geq k)) p_t^A.$$

Player S_R chooses $\{p_t^R\}_{t=2}^T \in \prod_{t=2}^T [\underline{\theta}_t, \bar{\theta}_t]$ to maximize

$$\sum_{t=2}^T \delta^{t-2} (1 - F_t(p_t^R | \theta_1 \leq k)) p_t^R.$$

The rest of proof is virtually the same as in [Appendix A.6](#) and thus omitted. \square

Step 2: We now consider the original subgame and show the following by induction.

Lemma 10. *The subgame after the seller chooses not to commit has a PBE- \star .*

Proof. Base case: When $T = 1$, the claim holds trivially.

Inductive step: Using the objects in [Lemma 9](#) (see [Appendix A.6](#) for the notation), we construct the strategies for the first two periods. Let the seller offer p^* in period 1. At any history, after the seller offers some price p in period 1, the buyer responds according to the threshold rule with threshold k^p . Then following histories A, R , the seller offers subsequent prices according to the distributional strategies η_A^p, η_R^p respectively; let her beliefs on $(\theta_1, \dots, \theta_T)$ be $F(\cdot | \theta_1 \geq k^p), F(\cdot | \theta_1 \leq k^p)$ respectively. The buyer accepts the offer in period t whenever the price in period t is weakly below θ_t for all $2 \leq t \leq T$.

To complete the construction, consider the off-path history where the seller chooses not to commit in period 2. Note that the seller's posterior belief about $(\theta_2, \dots, \theta_T)$ satisfies [Assumption A5](#) after any history (p, x_1) (since the buyer's first-period strategy is a threshold rule). Thus, by the inductive hypothesis, there exists a PBE- \star for the $(T - 1)$ -period game in which the buyer's type process is the posterior process $(\tilde{\theta}_2, \dots, \tilde{\theta}_T)$ (conditional

on $\theta_1 \geq k^p$ or $\theta_1 \leq k^p$), after any history (p, x_1) and the seller choosing not to commit. Use this PBE- \star to complete the construction of the strategies for the seller and buyer.

Given the seller's strategy, all types of the buyer are playing optimally by construction. By part (ii) of [Theorem 3](#) applying to the game starting from period 2, we know that after any history (p, x_1) the seller weakly prefers committing to some $\{p_2, \dots, p_T\}$. Then by the construction of the objects in [Lemma 9](#), the seller has no profitable deviation and has the required off-path beliefs. This concludes the inductive step. \square

Appendix B D1 Criterion

In this appendix, we provide a notion of D1 criterion in the spirit of [Banks and Sobel \(1987\)](#) for our game. As noted in [Section 4.2](#), the game defined in [Section 3](#) is a subgame of the modified multiple-period game defined in [Section 4.2](#). Thus, for notational convenience, we define all our notions in the multiple-period model (and allow the transition kernels of the type process to depend on the consumption in the previous period to accommodate [Section 4.1](#)). We also maintain the assumptions in [Section 4.1](#) for $T = 2$ (which nest those in [Section 3](#)), and the assumptions in [Section 4.2](#) for $T > 2$. Recall that we use PBE, PBE-D, PBE- \star to refer to the standard notion of PBE, PBE with D1 refinement (to be defined), and PBE with the particular class of off-path beliefs as in [Definition 1](#).

B.1 Definition of D1 Criterion

Fix any PBE. In any period t and at any buyer's history \hat{h}_t , a price p_t is posted by the seller. The buyer of type $\theta^t := (\theta_1, \dots, \theta_t)$ takes an action $x_t \in \{0, 1\}$. The seller observes this action, updates her belief, and moves to the continuation game. Let $U_t^*(\theta^t, \hat{h}_t)$ be type θ^t 's continuation utility in this PBE (including the payoff from this period). We suppress the dependency on \hat{h}_t whenever it is clear. For a given (potentially mixed) strategy α of the seller over the future plays in the continuation game, let $U_t(\theta^t, \alpha, \hat{h}_t)$ be the continuation utility of type θ^t best responding to α . As explained in [Lemma 5](#), the dependency on θ^t reduces to the dependency on θ_t . The seller also only makes inferences about θ_t from x_t .

Now fix any \hat{h}_t at which the PBE prescribes all types of the buyer to take the same action. Let μ_t be the prior belief on θ_t prescribed by the PBE for the seller before observing the buyer's action. Let x_t be the off-path action that the buyer could take. Let $BR(\Theta'_t, \hat{h}_t, x_t)$ denote the set of possible future plays α that are best responses to some buyer's strategy, under some belief ν_t on θ_t with the restriction that ν_t is only supported on $\Theta'_t \subseteq \text{supp}(\mu_t)$.

Now let

$$D(\theta_t, \Theta'_t, x_t) := \left\{ \alpha \in BR(\Theta'_t, x_t) \text{ s.t. } U_t^*(\theta_t) < U_t(\theta_t, \alpha, x_t) \right\}$$

denote the set of seller's possible strategies for the future play under some belief concentrated on Θ'_t that makes type θ_t strictly prefer x_t to the equilibrium play. Similarly, define

$$D^0(\theta_t, \Theta'_t, x_t) := \left\{ \alpha \in BR(\Theta'_t, x_t) \text{ s.t. } U_t^*(\theta_t) = U_t(\theta_t, \alpha, x_t) \right\}.$$

Let $\hat{\Theta}_t = \text{supp}(\mu_t)$. A type θ_t is deleted if there is another type $\theta'_t \in \hat{\Theta}_t$ such that

$$\left[D(\theta_t, \hat{\Theta}_t, x_t) \cup D^0(\theta_t, \hat{\Theta}_t, x_t) \right] \subset D(\theta'_t, \hat{\Theta}_t, x_t).$$

Let $\Theta_t^*(x_t)$ be the set of surviving types until this process stops. We say that the PBE does not survive the D1 criterion if there exists some $\theta_t \in \Theta_t^*(x_t)$ such that

$$U_t^*(\theta_t) < \inf_{\alpha \in \mathcal{A}(\Theta_t^*(x_t), x_t)} U_t(\theta_t, \alpha, x_t)$$

where $\mathcal{A}(\Theta_t^*(x_t), x_t)$ is the set of seller's strategies for the future play that belong to a PBE-D of the continuation game with some prior belief about θ_t supported only on $\Theta_t^*(x_t)$. Note that this definition is recursive. In the last period, after the buyer's choice, the seller no longer makes any decision, and so this set is \emptyset . In period $T - 1$, this set is the set of seller's strategies that belong to a standard PBE of the one-period continuation game.

In a signaling game, the above definition reduces to the usual definition, because there is no continuation game. In particular, the set of mixed best responses $BR(\cdot)$ does not depend on what the sender does subsequently; the set of receiver's continuation equilibrium responses $\mathcal{A}(\cdot)$ reduces to the set of mixed best responses.

B.2 PBE-D and PBE-*

In this section, we show that the PBE-D in our setting are essentially PBE-*. This provides a foundation for the particular class of off-path beliefs we focus on and confirms the intuition that the D1 criterion effectively requires the seller to believe the buyer is of the type that would benefit the most upon an unexpected deviation.

Lemma 11. *At any history \hat{h}_t , the set of D1 surviving types after an off-path rejection and acceptance are given by $\Theta_t^*(0) = \{\underline{\theta}_t(\mu_t)\}$, $\Theta_t^*(1) = \{\bar{\theta}_t(\mu_t)\}$, where $\underline{\theta}_t(\mu_t) = \min\{\text{supp}(\mu_t)\}$, $\bar{\theta}_t(\mu_t) = \max\{\text{supp}(\mu_t)\}$ and μ_t is the seller's belief on θ_t at history h_t .*

Proof. We prove the case for an unexpected off-path rejection. The other case follows by

a symmetric argument. Note that the mixed best response set $BR(\cdot)$ for the seller contains all feasible future strategies since in the definition we allow for the buyer to use arbitrary strategy for the continuation play.¹⁶ For any seller's strategy α fixed, by [Lemma 5](#) (and [Lemma 3](#) for the case of $T = 2$), we have that

$$H(\theta_t, \alpha) := U_t^*(\theta_t) - U_t(\theta_t, \alpha, 0)$$

is a strictly increasing and continuous function of θ_t . This implies that for any $\hat{\theta}_t$ and any $\theta'_t < \theta_t \in \hat{\theta}_t$,

$$\left[D(\theta_t, \hat{\theta}_t, 0) \cup D^0(\theta_t, \hat{\theta}_t, 0) \right] \subseteq D(\theta'_t, \hat{\theta}_t, 0).$$

Our hope is to delete the types from the top to the bottom. This requires a strict set inclusion instead of the weak one. Consider the particular class of strategy prescribing the sale of all future items at a constant price \hat{p} . Note that $H(\theta_t, \cdot)$ is a continuous function of \hat{p} and for high enough $H(\theta_t, \hat{p}) = U_t^*(\theta_t) \geq 0$. Moreover, for any \hat{p} low enough (potentially negative), $H(\theta_t, \hat{p}) < 0$. Therefore, there exists $\hat{p}^*(\theta_t)$ such that $H(\theta_t, \hat{p}^*(\theta_t)) = 0$. This strategy is in the set $D^0(\theta_t, \hat{\theta}_t, 0)$ for any θ_t and $\hat{\theta}_t$. Fix any θ_t and any $\theta'_t < \theta_t$. Consider the seller's strategy of selling all future items at price $\hat{p}^*(\theta_t) + \epsilon$ for $\epsilon > 0$ small enough. This strategy is not in $D(\theta_t, \hat{\theta}_t, 0) \cup D^0(\theta_t, \hat{\theta}_t, 0)$ by construction, but this strategy would be in $D(\theta'_t, \hat{\theta}_t, 0)$ for ϵ small enough. This is because $H(\cdot, \cdot)$ is strictly increasing in the first argument and continuous in the second. This argument then shows that for any $\hat{\theta}_t$ and any $\theta'_t < \theta_t \in \hat{\theta}_t$,

$$\left[D(\theta_t, \hat{\theta}_t, 0) \cup D^0(\theta_t, \hat{\theta}_t, 0) \right] \subset D(\theta'_t, \hat{\theta}_t, 0).$$

Iterating this gives $\Theta_t^*(0) = \{\underline{\theta}_t(\mu_t)\}$. □

Lemma 12. *Every PBE- \star is a PBE-D.*

Proof. We prove this by induction on T .

Base case: When $T = 1$, these two concepts are both equivalent to standard PBE.

Inductive step: Suppose for any $\tilde{T} \leq T - 1$ periods of our game, every PBE- \star is a PBE-D. Consider a T -period game and fix a PBE- \star . Fix any history \hat{h}_t where all types are prescribed to accept. Consider the off-path deviation $x_t = 0$. Note that by [Lemma 11](#), we have $\Theta_t^*(0) = \underline{\theta}_t(\mu_t)$. By the definition of PBE,

$$U_t^*(\underline{\theta}_t(\mu_t)) \geq U_t(\underline{\theta}_t(\mu_t), \tilde{\sigma}^*, 0)$$

where $\tilde{\sigma}^*$ is the seller's equilibrium continuation play. By the definition of PBE- \star , $\tilde{\sigma}^*$ is

¹⁶If the buyer rejects all future offers, then any price path of the seller is in the $BR(\cdot)$ set.

part of a PBE- \star for the continuation game under the belief concentrated on $\underline{\theta}_t(\mu_t)$. Since the continuation game has $\leq T - 1$ periods, by the inductive hypothesis, $\tilde{\sigma}^* \in \mathcal{A}(\underline{\theta}_t(\mu_t), 0)$. Therefore,

$$U_t^*(\underline{\theta}_t(\mu_t)) \geq U_t(\underline{\theta}_t(\mu_t), \tilde{\sigma}^*, 0) \geq \inf_{\alpha \in \mathcal{A}(\underline{\theta}_t(\mu_t), 0)} U_t(\underline{\theta}_t(\mu_t), \alpha, 0)$$

which shows that at any such history the PBE- \star passes the D1 test. The same argument holds for an unexpected acceptance. This proves the inductive step. \square

Lemma 13. *For every PBE-D, there exists an outcome-equivalent PBE- \star .*

Proof. We prove this by induction on T .

Base case: When $T = 1$, these two concepts are both equivalent to standard PBE.

Inductive step: Suppose for any $\tilde{T} \leq T - 1$ periods of our game, for every PBE-D, there exists an outcome-equivalent PBE- \star .

Consider a T -period game and fix a PBE-D. Fix any history \hat{h}_t where all types are prescribed to accept. Consider the off-path deviation $x_t = 0$. Note that by [Lemma 11](#), we have $\Theta_t^*(0) = \underline{\theta}_t(\mu_t)$. By the definition of PBE-D,

$$U_t^*(\underline{\theta}_t(\mu_t)) \geq \inf_{\alpha \in \mathcal{A}(\underline{\theta}_t(\mu_t), 0)} U_t(\underline{\theta}_t(\mu_t), \alpha, 0).$$

Since the continuation game has $\leq T - 1$ periods, by the inductive hypothesis,

$$\inf_{\alpha \in \mathcal{A}(\underline{\theta}_t(\mu_t), 0)} U_t(\underline{\theta}_t(\mu_t), \alpha, 0) = \inf_{\beta \in \mathcal{B}(\underline{\theta}_t(\mu_t), 0)} U_t(\underline{\theta}_t(\mu_t), \beta, 0)$$

where $\mathcal{B}(\underline{\theta}_t(\mu_t), 0)$ is the collection of seller's strategies that belong to a PBE- \star of the continuation game that is outcome-equivalent to a PBE-D in $\mathcal{A}(\underline{\theta}_t(\mu_t), 0)$.

If $T = 2$, then in the continuation game, PBE-D, PBE- \star , and PBE exist and coincide. They are specified by the distributions of prices that the seller posts. Let p^* be the highest price the seller posts across the equilibria. Replace the continuation strategy profile given in the original PBE-D by letting the seller post p^* and buyer accept if $\theta_2 \geq p^*$, with the seller's off-path belief assigned to be concentrated on $\underline{\theta}_t(\mu_t)$.

Now suppose $T > 2$. Note that the induction proof of part (ii) of [Theorem 3](#) can be done for PBE- \star without referring to PBE-D. This implies that the PBE- \star corresponding to $\mathcal{B}(\underline{\theta}_t(\mu_t), 0)$ are all outcome-equivalent. The existence of such a PBE- \star follows from the proof in [Appendix A.7](#) (and thus $\mathcal{B}(\underline{\theta}_t(\mu_t), 0)$ and $\mathcal{A}(\underline{\theta}_t(\mu_t), 0)$ are non-empty). Now select any such PBE- \star of the continuation game. Replace the continuation strategy profile given in the original PBE-D by this PBE- \star , with the seller's off-path belief assigned to be concentrated on $\underline{\theta}_t(\mu_t)$.

In both cases ($T = 2$ or $T > 2$), we claim that the resulting strategy profile for the whole game is still a PBE. To see this, note that by construction in both cases the replacement leaves $\underline{\theta}_t(\mu_t)$ a payoff of $\inf_{\alpha \in \mathcal{A}(\underline{\theta}_t(\mu_t), 0)} U_t(\underline{\theta}_t(\mu_t), \alpha, 0) \leq U_t^*(\underline{\theta}_t(\mu_t))$. So it remains optimal for type $\underline{\theta}_t(\mu_t)$ to play the prescribed strategy of acceptance. By the threshold-rule results (Lemma 3 and Lemma 5), it is then optimal for all types above $\underline{\theta}_t(\mu_t)$ to continue accepting the offer. Now consider one layer above where the seller is choosing what price to offer in period t . Because this replacement only happens at an off-path history and the buyer uses the same threshold rule as before for this period, it is optimal for the seller to follow the originally prescribed strategy as well. It then follows that after this replacement the strategy profile is still a PBE of the whole game, and has the same equilibrium outcome as the original PBE-D. The same procedure also works for any history \hat{h}_t at which all types of the buyer are prescribed to reject.

Now for each period $t \in \{1, \dots, T\}$, for each \hat{h}_t , iteratively apply the above procedure whenever (i) all types of the buyer are prescribed to accept or reject at \hat{h}_t and (ii) the assigned off-path belief is not already concentrated on $\underline{\theta}_t(\mu_t)$ or $\bar{\theta}_t(\mu_t)$. It is evident that when the process stops, we have found a PBE- \star that is outcome-equivalent to the given PBE-D. This concludes the inductive step. \square

References

- ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): “The Economics of Privacy,” *Journal of Economic Literature*, 54(2), 442–492.
- ACQUISTI, A. AND H. R. VARIAN (2005): “Conditioning Prices on Purchase History,” *Marketing Science*, 24(3), 367–381.
- AVIV, Y. AND A. PAZGAL (2008): “Optimal Pricing of Seasonal Products in the Presence of Forward-looking Consumers,” *Manufacturing & Service Operations Management*, 10(3), 339–359.
- BANKS, J. S. AND J. SOBEL (1987): “Equilibrium Selection in Signaling Games,” *Econometrica*, 55(3), 647–661.
- BARON, D. P. AND D. BESANKO (1984): “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy*, 1(3), 267–302.
- BATTAGLINI, M. (2005): “Long-term Contracting with Markovian Consumers,” *American Economic Review*, 95(3), 637–658.

- BERGEMANN, D. AND J. VÄLIMÄKI (2019): “Dynamic Mechanism Design: An Introduction,” *Journal of Economic Literature*, 57(2), 235–274.
- BONATTI, A. AND G. CISTERNAS (2020): “Consumer Scores and Price Discrimination,” *Review of Economic Studies*, 87(2), 750–791.
- CACHON, G. P. AND P. FELDMAN (2015): “Price Commitments with Strategic Consumers: Why It Can Be Optimal to Discount More Frequently . . . Than Optimal,” *Manufacturing & Service Operations Management*, 17(3), 399–410.
- COURTY, P. AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67(4), 697–717.
- CRÉMER, J. (1984): “On the Economics of Repeat Buying,” *RAND Journal of Economics*, 15(3), 396–403.
- DEB, R. AND M. SAID (2015): “Dynamic Screening with Limited Commitment,” *Journal of Economic Theory*, 159, 891–928.
- DEVANUR, N. R., Y. PERES, AND B. SIVAN (2019): “Perfect Bayesian Equilibria in Repeated Sales,” *Games and Economic Behavior*, 118, 570–588.
- DOVAL, L. AND V. SKRETA (2019): “Optimal Mechanism for the Sale of a Durable Good,” *SSRN Electronic Journal*.
- (2020): “Mechanism Design with Limited Commitment,” *SSRN Electronic Journal*.
- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2018): “Auctions versus Posted Prices in Online Markets,” *Journal of Political Economy*, 126(1), 178–215.
- ELMAGHRABY, W. AND P. KESKINOCAK (2003): “Dynamic Pricing in the Presence of Inventory Considerations: Research Overview, Current Practices, and Future Directions,” *Management Science*, 49(10), 1287–1309.
- ESO, P. AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705–731.
- FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): “Planning under Incomplete Information and the Ratchet Effect,” *Review of Economic Studies*, 52(2), 173–191.
- FUDENBERG, D. AND J. M. VILLAS-BOAS (2006): “Behavior-Based Price Discrimination and Customer Recognition,” *Handbook on Economics and Information Systems*, 377 – 436.

- (2012): “Price Discrimination in the Digital Economy,” *Oxford Handbook of the Digital Economy*, 254–272.
- GERARDI, D. AND L. MAESTRI (2020): “Dynamic Contracting with Limited Commitment and the Ratchet Effect,” *Theoretical Economics*, 15(2), 583–623.
- GLICKSBERG, I. L. (1952): “A Further Generalization of the Kakutani Fixed Theorem, with Application to Nash Equilibrium Points,” *Proceedings of the American Mathematical Society*, 3(1), 170–170.
- GUPTA, R. AND C. PATHAK (2014): “A Machine Learning Framework for Predicting Purchase by Online Customers Based on Dynamic Pricing,” *Procedia Computer Science*, 36, 599–605.
- HART, O. D. AND J. TIROLE (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55(4), 509.
- JING, B. (2011): “Pricing Experience Goods: The Effects of Customer Recognition and Commitment,” *Journal of Economics & Management Strategy*, 20(2), 451–473.
- KEHOE, P. J., B. J. LARSEN, AND E. PASTORINO (2020): “Dynamic Competition in the Era of Big Data,” *Working Paper, Stanford University and Federal Reserve Bank of Minneapolis*.
- KENNAN, J. (2001): “Repeated Bargaining with Persistent Private Information,” *Review of Economic Studies*, 68(4), 719–755.
- LAFFONT, J.-J. AND J. TIROLE (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 56(5), 1153.
- LIU, Q., K. MIERENDORFF, X. SHI, AND W. ZHONG (2019): “Auctions with Limited Commitment,” *American Economic Review*, 109(3), 876–910.
- MILGROM, P. R. AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089.
- (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10(4), 619–632.
- PAVAN, A. (2017): “Dynamic Mechanism Design: Robustness and Endogenous Types,” in *Advances in Economics and Econometrics*, 1–62.

- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic mechanism design: A Myersonian approach,” *Econometrica*, 82, 601–653.
- SKRETA, V. (2006): “Sequentially Optimal Mechanisms,” *Review of Economic Studies*, 73(4), 1085–1111.
- STOKEY, N. L. (1979): “Intertemporal Price Discrimination,” *Quarterly Journal of Economics*, 93(3), 355–371.
- STRULOVICI, B. (2017): “Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts,” *Econometrica*, 85(2), 585–616.
- SU, X. (2007): “Intertemporal Pricing with Strategic Customer Behavior,” *Management Science*, 53(5), 726–741.
- TAYLOR, C. R. (2004): “Consumer Privacy and the Market for Customer Information,” *RAND Journal of Economics*, 35(4), 631–650.
- THOMPSON, F. M. AND T. CHMURA (2015): “Loyalty Programs in Emerging and Developed Markets: The Impact of Cultural Values on Loyalty Program Choice,” *Journal of International Marketing*, 23(3), 87–103.
- TOPKIS, D. M. (1998): *Supermodularity and Complementarity*, Princeton University Press.
- VILLAS-BOAS, J. M. (2004): “Price Cycles in Markets with Customer Recognition,” *RAND Journal of Economics*, 35(3), 486.