Multilateral Bargaining with Information Spillovers*

Vladimir Asriyan William Fuchs Caio Lorecchio

October 30, 2026

Abstract

We study dynamic monopoly pricing with multiple buyers whose privately known valuations are correlated, so one buyer's purchase reveals information about others. Without seller commitment, these information spillovers alter future pricing incentives and feed back into current demand. They generate strategic complementarities in buyers' purchase decisions, making demand nonmonotone and giving rise to equilibrium multiplicity. Information spillovers can strengthen or weaken the seller's effective commitment power, implying that seller revenue, buyer surplus, and welfare can be nonmonotone in spillover strength. We characterize when these spillovers encourage data investment and experimentation, as well as facilitate coordination that mitigates hold-up.

JEL Classification: D42, D82, D83.

Keywords: Dynamic monopoly pricing, information spillovers, strategic complementarities, equilibrium multiplicity, commitment.

*Asriyan: CREi, ICREA, BSE and CEPR (vasriyan@crei.cat). Fuchs: University of Texas at Austin and CEPR (william.fuchs@mccombs.utexas.edu). Lorecchio: Universitat de Barcelona, BEAT and CREB (paeslemelorecchio@ub.edu). The order of authors is alphabetical. We thank Berardino Cesi, Laurenz Marstaller, Goonj Mohan, Dmitry Orlov, Andrzej Skrzypacz, Tsz-Ning Wong and seminar participants at CREi-UPF, Federal University of Minas Gerais, Universitat de Barcelona, Universitat de les Illes Balears, University of São Paulo, University of Duisburg-Essen, University of Mannheim, UB Microeconomics Workshop, Catalan Economic Society Conference, Jornadas de Economía Industrial, Stony Brook Conference, International Industrial Organization Conference, Annual Meetings of the Armenian Economic Association, Durham Economic Theory Conference, Meeting of the Society for the Advancement of Economic Theory, International Conference on Public Economic Theory and International Workshop on Game Theory and Economic Applications for helpful comments and suggestions. Asriyan acknowledges financial support from the European Research Council (MACFRIC, grant 101229916) and the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2024-001476-S). Fuchs acknowledges financial support from grant PGC2018-096159-B-I00, financed by MCIN/AEI, and from the Comunidad de Madrid through grants EPUC3M11 (V PRICIT) and H2019-HUM-5891. Lorecchio acknowledges financial support from grant SGR2021-0082, financed by GENCAT-AGAUR, and grants PID2021-128237OB-I00, PID2023-148696OA-I00 and PID2024-158983OB-100, financed by MCIN/AEI.

1 Introduction

In many markets, buyers' valuations are correlated, allowing firms to learn about one buyer from the behavior of others. Online retailers rely on market-wide transaction data to tailor prices, while real-estate platforms such as Zillow and Redfin make recent neighborhood transactions publicly observable.[1] With correlated valuations, an observed sale reveals information about the demand the seller still faces and therefore affects his incentive to maintain or adjust future prices. We refer to this as an *information spillover*.

We study how information spillovers reshape dynamic monopoly pricing when sellers cannot commit to future prices (Coase 1972). We consider a two-period monopolist selling to two privately informed buyers with positively correlated valuations. First-period trades are publicly observed before the seller sets second-period prices. The seller cannot commit to future prices and, in our baseline, cannot distinguish buyers, so he posts a single price each period. Absent correlation, the model reduces to the textbook dynamic monopoly problem, where the seller screens each buyer separately (Sobel and Takahashi 1983; Fudenberg and Tirole 1983; Fuchs and Skrzypacz 2013). Correlation creates an informational link across buyers: a purchase by one buyer changes what the seller learns about the other buyer, coupling the seller's pricing decisions across buyers. The information conveyed by trade is *endogenous*: how much an early purchase reveals depends on how aggressively high-valuation buyers trade, which in turn depends on the continuation prices they expect.

Two opposing forces are at work. On the one hand, the absence of early trade is bad news: because high-valuation buyers purchase earlier than low-valuation buyers, it makes the seller more pessimistic about remaining demand and more willing to cut future prices—the familiar skimming force, which encourages delay. On the other hand, an early purchase is good news: with correlated valuations, it makes the seller more optimistic about remaining demand and, therefore, more willing to keep future prices high. Crucially, the same first-period trading probability governs the strength of both forces.

This second force fundamentally changes the monopolist's pricing problem by generating strategic complementarities in buyers' purchase decisions. As a result, the monopolist's demand correspondence can become nonmonotone in price. Because a high-valuation buyer expects a high continuation price precisely when the seller is optimistic about remaining demand—that is, when other high-valuation buyers trade early—her willingness to buy early rises with expected early trade by others. A higher first-period price can therefore be sustained by a higher probability of early trade: buyers accept the higher price because they expect waiting to be less attractive. This feedback between prices, beliefs, and trading behavior lies at the heart of our results and is absent from the standard dynamic monopoly problem with independent values.

Our first main result characterizes equilibria as a function of the strength of information spillovers (Theorem 1). There is a correlation threshold below which the equilibrium is unique and takes one of two familiar forms: a *high-price* equilibrium, in which only high-valuation buyers trade over time, or a *low-price* equilibrium, in which the seller cuts his second-period price to sell to everyone. Above the threshold, information spillovers are strong enough that the feedback between prices and beliefs sustains a continuum of

[1] In financial markets, post-trade transparency rules, such as TRACE for the U.S. corporate bond markets, mandate publishing completed transactions, precisely so that participants can learn from them.

contingent-price equilibria in which future prices depend on the information revealed by early trades and which may coexist with either the high-price or low-price equilibrium. Correlation therefore transforms a standard screening problem into one in which information links buyers' incentives and generates equilibrium multiplicity. Moreover, the multiplicity is economically consequential, as equilibria differ sharply in prices, trade, and the allocation of surplus.

Our second set of results concerns how information spillovers shape the seller's effective commitment power and the welfare consequences of dynamic pricing. Under commitment, spillovers never hurt the seller (Proposition 1): correlation reduces information rents, and as correlation approaches one the seller captures the surplus from an allocation in which high-valuation buyers purchase immediately and low-valuation buyers purchase in the second period. Without commitment, however, spillovers may either strengthen or weaken the seller's effective commitment power. They strengthen it when favorable early sales support high future prices, but weaken it when weak early sales induce future price cuts that buyers anticipate. Above the multiplicity threshold, the highest equilibrium revenue increases with correlation and converges to the commitment benchmark, whereas the lowest equilibrium revenue remains bounded away from it (Proposition 5). Moreover, this possibility of approaching commitment is not a consequence of the two-period horizon: in an infinite-horizon extension, for any fixed discount factor below one, there exists a sequence of equilibria in which the seller's revenue approaches the commitment payoff as correlation approaches one (Proposition B.1). Buyer surplus and welfare, by contrast, change discontinuously as the equilibrium set changes (Proposition 6). Consequently, revenue, buyer surplus, and welfare can be nonmonotone in correlation and need not move together.

These forces have implications for data investments, experimentation, and market design. Richer consumer data lets firms segment buyers into finer, more homogeneous groups (Bergemann, Brooks and Morris 2015), which we represent in our model as a costly increase in the correlation of buyers' valuations. We show that a seller makes this investment only when it makes information spillovers sufficiently strong, and that it can either raise or lower buyer surplus and welfare depending on the equilibrium otherwise played (Corollary 1). This qualifies the presumption that richer consumer data is unambiguously harmful for consumers and identifies when restrictions on data use, such as the EU's General Data Protection Regulation and Digital Services Act, improve welfare. A similar logic applies to experimentation: by using an early offer to one buyer as an informational probe, the seller refines the subsequent offer to the other buyer, making experimentation profitable only when information spillovers are sufficiently strong (Corollary 2). A different implication of equilibrium multiplicity is that it can mitigate hold-up (Corollary 3). In a procurement interpretation of the model, sufficiently strong spillovers and sufficiently productive investments allow suppliers to coordinate on equilibria that preserve enough rents to support relationship-specific investment. This suggests a rationale for mechanisms that facilitate coordination among suppliers.

In sum, our results show that information revealed by trade is not merely a by-product of market interactions but a source of market power when sellers cannot commit. Endogenizing this information fundamentally reshapes dynamic monopoly pricing, with implications for prices, trade, and surplus allocation.

Related literature. We contribute to the literature on dynamic monopoly pricing with forward-looking buyers. Since Coase (1972), a central theme in this literature is that a seller who cannot commit to future

prices may be unable to extract rents from buyers. Formal analyses by Stokey (1979), Bulow (1982), Sobel and Takahashi (1983), Fudenberg and Tirole (1983), and Gul, Sonnenschein and Wilson (1986) show how the prospect of future price cuts constrains current prices. In these models, the seller's dynamic problem is one of screening buyers where each buyer's purchase reveals information only about that buyer's willingness to pay. We introduce a distinct channel: a buyer's purchase also reveals information about the valuations of other buyers. This cross-buyer learning changes the seller's continuation incentives and feeds back into current demand, generating strategic complementarities, nonmonotone demand, and equilibrium multiplicity that are absent from the standard settings with independent private values.

Related work studies revenue management and intertemporal pricing when buyers strategically time purchases. Board and Skrzypacz (2016) characterize the revenue-maximizing mechanism for a seller facing a random arrival of forward-looking buyers, and Libgober and Mu (2021) study intertemporal pricing when the seller is uncertain about how a buyer learns her valuation over time. These papers hold buyers' information fixed or treat its evolution as exogenous. In our setting, instead, information about a delaying buyer is generated endogenously by the purchasing decisions of other buyers.

The paper most closely related to ours is Drugov (2021), who studies dynamic bargaining with correlated private information across multiple bilateral negotiations. In his model, agreement in one relationship reveals information about another, leading principals to benchmark agents against one another and generating equilibrium multiplicity. We instead study a monopolist selling to multiple buyers with correlated valuations. Because the monopolist prices all buyers, he internalizes how current prices affect the information revealed by trade and, through it, future pricing incentives. This endogenous interaction between pricing and information revelation generates nonmonotone demand and reshapes the seller's effective commitment power.

Our mechanism is also related to Asriyan, Fuchs and Green (2017), which studies a dynamic market for lemons in which the trade of one asset reveals information about the value of another. As in that paper, sufficiently strong information spillovers generate strategic complementarities and equilibrium multiplicity. Unlike their competitive market setting with price-taking informed traders, our setting features a monopolist who strategically sets prices but cannot commit to future pricing. Spillovers therefore affect not only equilibrium trading behavior but also the seller's continuation pricing incentives.

Our work also relates to the literature on social learning. One strand studies dynamic pricing when earlier purchases convey information about a common payoff-relevant state to subsequent buyers (Bose et al. 2006; Arieli, Koren and Smorodinsky 2022; Laiho and Salmi 2026); another studies strategic delay, in which agents postpone actions to learn from others' behavior, generating strategic complementarities and multiple equilibria (Chamley and Gale 1994). Our setting features a different informational channel: buyers do not learn an unknown common state from previous actions, but a buyer's purchase changes the seller's beliefs about other buyers' valuations and thus continuation prices, and it is the anticipation of this pricing response that links buyers' purchasing decisions and generates strategic complementarities.

Our applications also relate to the literature on dynamic information provision and recommendation (Kremer, Mansour and Perry 2014; Che and Hörner 2018), and to work on information design and price discrimination (Roesler and Szentes 2017; Bonatti and Cisternas 2020; Bergemann, Brooks and Morris 2025; Kartik and Zhong 2026; Elliott et al. 2026). In these papers, an intermediary or designer controls the

information disclosed to buyers or firms. In our baseline setting, by contrast, the informational content of histories is determined endogenously by equilibrium trading decisions. Nevertheless, our applications to data investment and A/B testing can be viewed as constrained forms of information design: by changing the correlation of buyers' valuations or using some buyers as informational probes, the seller changes the informativeness of purchase histories and, in turn, buyers' intertemporal decisions.

Finally, we contribute to the literature on data, pricing, and market transparency. One strand studies price discrimination based on exogenous consumer characteristics or purchase histories, including personalized pricing (Thisse and Vives 1988) and behavior-based pricing that conditions offers on past purchases (Fudenberg and Villas-Boas 2006). Another strand examines the effects of transaction transparency on market outcomes (Fuchs, Öry and Skrzypacz 2016). Our focus is instead on how sellers' pricing incentives respond to information about remaining demand that is revealed endogenously through equilibrium trade. Consequently, policies that affect data use, market segmentation, or transaction transparency influence market outcomes not only by changing the scope for price discrimination, as emphasized in the existing literature, but also by altering sellers' continuation pricing incentives. Through this channel, they affect trade, welfare, and the seller's ability to extract surplus.

Outline. Section 2 presents the model and establishes the commitment benchmark. Section 3 characterizes the equilibrium set and presents our main results. Section 4 studies how information spillovers affect the seller's effective commitment power, welfare, and the allocation of surplus. Section 5 develops applications to data investment, A/B testing, and supplier coordination. Section 6 discusses extensions that demonstrate the robustness and broader applicability of our mechanism. Proofs and formal derivations are in the appendices.

2 The Model

Consider a monopolist selling a product to two buyers, indexed by $i \in \{A, B\}$, who must decide whether and when to make a purchase. Each buyer can purchase at most one unit and has at most two periods to do so. The seller posts a unit price in each period and has a constant marginal cost of production, normalized to zero. Buyer i can be one of two types, $\theta_i \in \{H, L\}$, and values the product at v_{θ_i}, with $v_H > v_L > 0$.

All agents discount future payoffs using a common factor $\delta \in (0, 1)$. Therefore, the payoff to buyer i of type θ_i from purchasing the good at price p_t in period $t \in \{1, 2\}$ is given by:

$$\delta^{t-1}(v_{\theta_i} - p_t). \tag{1}$$

Otherwise, if the buyer never trades, she obtains a payoff of zero.

The payoff to the seller from producing and selling the good at price p_t in period t is:

$$\delta^{t-1}p_t. \tag{2}$$

If no trade occurs, the seller likewise receives a payoff of zero.

While each buyer knows her own valuation, she is uninformed about the valuation of the other buyer. Importantly, the seller does not know the valuation of either buyer, giving rise to a problem of asymmetric

information between the seller and the buyers.

The central feature of our model is that buyers' valuations are positively correlated. As a result, the trading behavior of one buyer may provide the seller with information about the valuation of the other buyer. Let the unconditional distribution of buyer types be:

$$\mathbb{P}(\theta_i = H) = \pi_1, \tag{3}$$

for $i \in \{A, B\}$ with $\pi_1 \in (0, 1)$. The conditional distribution of θ_i given θ_j is:

$$\mathbb{P}(\theta_i = H | \theta_j = H) = \lambda, \tag{4}$$

for $i, j \in \{A, B\}$, $j \neq i$, and $\lambda \in [\pi_1, 1)$. The parameter λ captures the degree of correlation: when $\lambda = \pi_1$, valuations are uncorrelated; as $\lambda \to 1$, correlation becomes perfect.[2]

To streamline the analysis, we impose the following parametric assumption:

Assumption 1 (No pooling)**.** $\pi_1 v_H > v_L$.

Assumption 1 ensures that asymmetric information is severe enough, so that in a one-shot interaction the seller prefers not to trade with all buyer types. If the seller were to serve both types at the highest price acceptable to both, his payoff would be v_L. By contrast, trading only with high types yields an expected payoff $\pi_1 v_H$, which is strictly higher under the assumption. While this assumption simplifies the analysis, the mechanism generating strategic complementarities, nonmonotone demand, and equilibrium multiplicity does not rely on it (see Appendix B).

Throughout, we assume that the seller is unable to price discriminate buyers solely based on their identity:

Assumption 2 (Anonymity)**.** *The identity $i \in \{A, B\}$ of a buyer is not observable to the seller. Consequently, all buyers must be offered the same price in a given period.*

One justification for Assumption 2 is that the product market is centralized, allowing a buyer offered a lower price to resell the product to another buyer. Alternatively, the restriction may result from regulations aimed at leveling the playing field among consumers.[3] We return to this interpretation in Section 5, where we relax Assumption 2 and examine the implications of our theory for policies aimed at consumer data protection.

Strategies and Information Sets. A strategy for the seller is a deterministic price in the first period and a mapping from his information set to a probability distribution over prices in the second period.[4] In the first period, his information set is empty. In the second period, it includes the trading history in the first period, i.e., the first-period offer and how many buyers have accepted it.

A strategy for buyer i maps her information set to a probability of buying the product. In the first period, buyer i's information set includes her valuation and the first-period offer posted by the seller. In the second

[2] A correlation λ also implies that $\mathbb{P}(\theta_i = H | \theta_j = L) = \frac{\pi_1(1-\lambda)}{1-\pi_1}$. We assume $\lambda < 1$ to avoid the need to specify off-equilibrium beliefs. Nevertheless, equilibria with $\lambda = 1$ can be shown to be the limits of equilibria with $\lambda < 1$ as λ approaches one.

[3] For example, the EU's Geo-blocking Regulation prohibits discrimination based on customers' nationality, place of residence, or place of establishment. See https://eur-lex.europa.eu/content/news/geo-blocking-regulation-enters-into-force.html.

[4] While we could also allow the seller to mix over prices in the first period, our main results do not depend on this restriction; relaxing it would only make the welfare set convex (see Section 4).

period, her information set also includes the second-period offer posted by the seller and whether buyer j traded in the first period.

Equilibrium. Our equilibrium concept is perfect Bayesian. This has three implications.

1. *Buyer Optimality*: each buyer's acceptance rule must maximize her expected payoff at every information set, taking as given the other buyer's acceptance rule and the seller's pricing strategy.

2. *Seller Optimality*: the seller's pricing strategy must maximize his expected profits following each trading history, given his beliefs and the buyers' acceptance rules.

3. *Belief Consistency*: the seller's beliefs must be consistent with Bayes' rule whenever possible.

It is useful to clarify how our equilibrium concept also puts restrictions on play following an unexpected deviation by the seller in the first period. To this end, suppose that in equilibrium the seller was to make an offer p_1 in the first period, but an offer $\tilde{p}_1$ was observed instead. Since the seller is uninformed, such a deviation does not contain any information. Thus, we require that the strategies that follow $\tilde{p}_1$ are consistent with the associated conditional beliefs, and that these beliefs are updated by Bayes' rule.

We note that our equilibrium concept presumes a lack of commitment on the part of the seller, since he must act optimally in the second period following every history. We will, however, consider a benchmark with seller commitment, where we drop this requirement.

2.1 Preliminaries

It is convenient to highlight some basic properties that all equilibria must satisfy.

Continuation values. We assume without loss of generality that the support of prices offered by the seller in the second period is $\{v_L, v_H\}$. Any price above v_H (below v_L) is rejected (accepted) for sure by both buyers. In addition, any price in (v_L, v_H) would only be accepted by the H-type, so the seller could profitably deviate by raising it to v_H. Thus, in any equilibrium, if a buyer delays trade to the second period, she is later offered either v_H or v_L, and accepts whenever her valuation is at least that offer.[5]

Given a trading history z in the first period, let $\phi(z)$ denote the probability that the seller offers v_H in the second period. Denote by $\pi_2(z)$ the seller's posterior belief that the delaying buyer has a high valuation, i.e., v_H. The seller's expected second-period revenue from trading with this buyer after the history z is:

$$\phi(z)\pi_2(z)v_H + (1 - \phi(z))v_L. \tag{5}$$

That is, with probability $\phi(z)$, the seller offers v_H, which is accepted by the buyer with probability $\pi_2(z)$, i.e., when she is an H-type. Instead, with probability $1 - \phi(z)$, the seller offers v_L, which is accepted for sure by both buyer types.

If the seller were able to commit to second-period offers, he could choose any $\phi(z) \in [0, 1]$. Lack of commitment, however, requires that $\phi(z)$ be optimal for the seller *ex post*. Define $\pi^\star := \frac{v_L}{v_H}$, which is smaller than π_1 by Assumption 1. In particular, from (5):

[5] If a buyer accepted with probability less than one when indifferent, the seller could deviate by slightly lowering the offer.

$$\phi(z) \begin{cases} = 1 & \text{if } \pi_2(z) > \pi^\star \\ \in [0,1] & \text{if } \pi_2(z) = \pi^\star \\ = 0 & \text{if } \pi_2(z) < \pi^\star \end{cases}. \tag{6}$$

Given the seller's second-period strategy and the buyers' optimal responses, the second-period surplus $S_\theta(z)$ of a buyer of type θ after history z is:

$$S_H(z) = [1 - \phi(z)](v_H - v_L) \text{ and } S_L(z) = 0. \tag{7}$$

As a result, when deciding whether to accept an offer in the first period, the continuation value of buyer i of type θ from delaying trade to the second period is:

$$Q_H^i = \delta E_H^i \left[S_H(z)\right] \text{ and } Q_L^i = 0, \tag{8}$$

where $E_H^i[\cdot]$ indicates that the expectation is taken with respect to the information set of H-type buyer i. In principle, if buyers' strategies were asymmetric in the first period, $E_H^i[\cdot]$ could be buyer-specific. However, given buyer symmetry and market anonymity, we can show that all equilibria are symmetric (see Lemma B.1 in Appendix B). Thus, we drop the superscript i from expectations and continuation values in what follows.

First-period trade. The well-known skimming property holds in our setting.[6] That is, any offer that is acceptable to the L-type buyer must be accepted with probability one by the H-type buyer. As a result:

Lemma 1. *In any equilibrium, the seller's first-period offer is weakly above $v_H - \delta(v_H - v_L)$, which is never accepted by the L-type and is accepted by the H-type with some probability $\sigma \geq 0$.*

Suppose to the contrary that both buyer types accept the seller's first-period offer. Then it must be weakly below v_L. However, the seller could profitably deviate by offering $v_H - \delta(v_H - v_L) - \varepsilon > v_L$, which the H-type would accept with certainty, since v_L is the lowest offer the buyer expects in the second period. This deviation is profitable because Assumption 1 and a small enough ε ensure that $\pi_1\big[v_H - \delta(v_H - v_L) - \varepsilon\big] + (1 - \pi_1)\delta v_L > v_L$. Together with the skimming property, this implies that the L-type never trades in the first period. Moreover, at any first-period offer $p_1 < v_H - \delta(v_H - v_L)$, the H-type accepts with probability one regardless of continuation play, so the posterior after either history is zero and the seller's per-buyer revenue is $\pi_1 p_1 + (1 - \pi_1)\delta v_L$; a slightly higher offer is then strictly profitable. Hence, in any equilibrium, the first-period offer is weakly above $v_H - \delta(v_H - v_L)$.

Updating. Before making second-period offers, the seller updates his beliefs about the remaining buyers. Hence, given the first-period offer, there are two relevant histories to consider: either no buyer traded or one buyer traded in the first period. Given that in equilibrium only the H-type accepts an offer in the first period (Lemma 1), we refer to the event that no buyer traded as "bad news", and denote it by $z_b := (b, p_1)$. Analogously, we refer to the event that only one buyer traded as "good news", and denote it by $z_g := (g, p_1)$.

[6] See, for example, Lemma 10.1 in Fudenberg and Tirole (1991).

Given the first-period trading probability σ of the H-type buyer, let $\rho_\theta(z;\sigma) \equiv \mathbb{P}_\sigma(z|\theta)$ denote the θ-buyer's belief that history $z \in \{z_b, z_g\}$ is observed before the seller makes his second-period offer. Then:

$$\rho_\theta(z_g;\sigma) = \begin{cases} \lambda\sigma & \text{if } \theta = H \\ \frac{\pi_1(1-\lambda)}{1-\pi_1}\sigma & \text{if } \theta = L \end{cases}. \tag{9}$$

As a result, provided that $\sigma > 0$, the seller's second-period posterior belief, $\pi_2(z;\sigma)$, that a remaining buyer is an H-type, following the trading history z, can be expressed as:

$$\pi_2(z;\sigma) = \frac{(1-\sigma)\,\rho_H(z;\sigma)\pi_1}{(1-\sigma)\rho_H(z;\sigma)\pi_1 + \rho_L(z;\sigma)(1-\pi_1)}. \tag{10}$$

Otherwise, if $\sigma = 0$, we set $\pi_2(z;0) = \pi_1$.

Note that $\pi_2(z_g;\sigma) \geq \pi_2(z_b;\sigma)$, which justifies the labels "good news" and "bad news" respectively. Also, both $\pi_2(z_g;\sigma)$ and $\pi_2(z_b;\sigma)$ decrease in σ: the higher the trading probability of high-valuation buyers, the more pessimistic the seller's belief is following no trade.

2.2 Commitment Benchmark

Before turning to the equilibrium analysis, we consider a benchmark in which the seller can commit to second-period offers. The benchmark serves two purposes. It isolates the role of the seller's inability to commit, which is the only feature that distinguishes it from the main analysis, and it provides the upper bound against which we measure the seller's *effective* commitment power in Section 4.

Pricing plans. Before the first period, the seller publicly announces a pricing plan: a first-period price p_1 together with a rule that assigns a second-period price, possibly random, to each first-period trading history $z \in \{z_g, z_b\}$. Since restricting second-period prices to $\{v_L, v_H\}$ remains without loss, a plan is a triple $(p_1, \phi(z_g), \phi(z_b))$, where $\phi(z)$ is the probability of posting v_H after history z, as in Section 2.1. The plan is binding: once the history is realized, the seller posts the announced price and cannot revise it in light of his posterior $\pi_2(z)$. In terms of our equilibrium concept, the benchmark retains Buyer Optimality and Belief Consistency but requires Seller Optimality only at the announcement stage. The ex post constraint (6) no longer applies, and the seller may choose any $\phi(z) \in [0,1]$.

Everything else is as in the main analysis. The seller sells through posted prices, posts a single price to all remaining buyers in each period (Assumption 2), and can condition second-period prices only on the publicly observed trading history. He cannot elicit reports from buyers, make payments contingent on such reports, or allocate the good other than through take-it-or-leave-it offers. Commitment thus enlarges the set of second-period prices the seller can credibly promise, but not the set of instruments through which he trades.

Posted prices versus mechanisms. We deliberately do not endow the committed seller with arbitrary mechanisms. First, because the two settings share the same instruments, any difference between the benchmark and the main analysis is due to constraint (6) alone. Second, a full mechanism-design benchmark would be degenerate in our environment. Whenever $\lambda > \pi_1$, an H-type and an L-type hold different beliefs about

the other buyer's valuation, so as in Crémer and McLean (1988) a seller who could commit to an arbitrary mechanism would extract the entire surplus in the first period, however weak the correlation. Such a benchmark is silent about our central question: how the strength of information spillovers shapes outcomes. Third, posted prices are closer to how sellers operate in the markets we have in mind. Sellers sometimes commit to future prices, for instance through pre-announced price paths or price guarantees, whereas mechanisms that deliver full extraction have no natural counterpart in these markets.

Equilibrium selection. A pricing plan may induce multiple equilibria in the game among buyers.[7] In particular, when $\phi(z_g) > \phi(z_b)$, H-types' first-period purchase decisions are strategic complements: expecting the other buyer to purchase raises the likelihood of good news and hence of a high continuation price, making immediate purchase more attractive. For instance, under the plan $p_1 = v_H - \delta(1-\lambda)(v_H - v_L)$, $\phi(z_g) = 1$, and $\phi(z_b) = 0$, both immediate acceptance and joint delay by the H-types are equilibria of the buyers' game. We select, for each plan, the buyers' equilibrium that maximizes the seller's revenue, which gives commitment its best chance.[8] We stress that this selection is used only for the benchmark. In the main analysis, we impose no selection and characterize the entire equilibrium set; there, multiplicity is a substantive implication of information spillovers (Theorem 1).

A key feature of the benchmark is that information spillovers only help the seller: they allow him both to (i) realize gains from trade with low-valuation buyers and (ii) discipline delays by high-valuation buyers.

Proposition 1. *In the commitment benchmark, the seller's optimal pricing plan is generically unique.*[9] *Under it, the H-type buyer always trades in the first period, and there exists a threshold $\lambda^c \in (\pi_1, 1)$ such that:*

- ***High price**: if $\lambda < \lambda^c$, the seller posts v_H in both periods, irrespective of news, and the L-type never trades;*
- ***Contingent price**: if $\lambda > \lambda^c$, the seller posts $p_1 = v_H - \delta(1-\lambda)(v_H - v_L)$ in the first period and, in the second, v_H after good and v_L after bad news; the L-type trades if and only if there is bad news.*

Under the optimal plan, high types trade immediately, so only low-valuation buyers remain in the second period. Every second-period offer of v_H is therefore ex post suboptimal: a seller free to revise his price would cut it to v_L, as constraint (6) requires. The value of committing to v_H lies entirely in its effect on first-period demand: it removes the high type's option to wait for a discount. When valuations are uncorrelated ($\lambda = \pi_1$), news carries no information and, as in Stokey (1979) and Bulow (1982), the seller optimally commits to v_H in both periods and earns $\pi_1 v_H$. This remains optimal when correlation is positive but weak ($\lambda < \lambda^c$).

When correlation is strong ($\lambda > \lambda^c$), the seller does better with a news-contingent plan. Cutting the price to v_L after bad news realizes the gains from trade with low types. Keeping it at v_H after good news is what allows the seller to keep the first-period price high: an H-type who contemplates delaying expects

[7] By the argument of Lemma 1, which applies equally when the seller commits, no optimal plan induces first-period trade by L-types, so it is without loss to restrict attention to plans under which only H-types trade in the first period.

[8] The selection is not innocuous. If buyers instead coordinated on delay whenever delay is an equilibrium of their game, the committed seller could still induce H-types to accept immediately by lowering p_1 slightly below v_H and setting $\phi(z) = 1$ for all z, in which case accepting is optimal regardless of the other buyer's behavior, but potentially at a revenue loss.

[9] By generically unique, we mean that nonuniqueness arises only in knife-edge parameter configurations.

good news with probability λ, and hence a discount only with probability $1-\lambda$, so a first-period concession of $\delta(1-\lambda)(v_H - v_L)$ suffices to induce her to trade immediately. Contingent pricing is optimal when the expected discounted gains from trade with low types, $\delta v_L(1 - 2\pi_1 + \pi_1\lambda)$, exceed the concession to high types, $\pi_1\delta(1-\lambda)(v_H - v_L)$, which holds if and only if $\lambda > \lambda^c = 1 - \frac{(1-\pi_1)v_L}{\pi_1 v_H}$. As $\lambda \to 1$, the concession vanishes and revenue converges to $\pi_1 v_H + (1-\pi_1)\delta v_L$, the entire surplus subject to delayed trade with low types. This result echoes Crémer and McLean (1988): buyers' information about one another reduces their rents. Because the seller is restricted to posted prices, however, information is revealed only through trade, so full extraction is approached only in the limit and only with delay.

3 Equilibrium Characterization

We now turn to our main setting, in which the seller lacks commitment to second-period offers. We begin by considering a fictitious game in which the first-period offer is treated as exogenous. For any such offer, we compute the corresponding acceptance probability of the H-type buyers, which defines a *demand schedule* faced by the seller in the first period. A key feature of this schedule is that, owing to information spillovers, it is nonmonotone (Proposition 2). Consequently, multiple trading probabilities may be consistent with a subgame equilibrium following the offer—a property that lies at the heart of equilibrium multiplicity.

Building on the demand schedule, we show that every equilibrium belongs to one of three classes: *high-price*, *low-price*, or *contingent-price* (Proposition 3). In the first two, the seller's second-period offers are not contingent on news: he charges v_H in the high-price equilibrium, where high types trade early with interior probability, and v_L in the low-price equilibrium, where high types trade with certainty. In the contingent-price equilibrium, by contrast, second-period prices depend on news, and multiple trading probabilities can be sustained. Theorem 1 characterizes the existence of all three classes and shows that multiplicity requires sufficiently strong information spillovers, both across classes and within the contingent-price class, where a continuum of equilibria emerges.

3.1 The Demand Schedule

For any first-period offer $\widetilde{p}_1$, there is an associated trading probability $\sigma \equiv \sigma(\widetilde{p}_1)$ with which the H-types accept it. This probability defines the seller's second-period posterior belief $\pi_2(z;\sigma)$ following each history $z \in \{z_b, z_g\}$, where $z_b = (b, \widetilde{p}_1)$ and $z_g = (g, \widetilde{p}_1)$ (see (10)). We now characterize the demand schedule and illustrate its key features in Figure 1.[10]

First, for any first-period offer above v_H, there will be no trade in the first period. At price v_H, the H-types are indifferent between trading and not trading, provided they expect a price v_H again in the second period with certainty. This belief is consistent with equilibrium play only if the seller's posterior belief is sufficiently high, i.e., $\pi_2(z;\sigma) \geq \pi^\star$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \leq \underline{\sigma}$ for $\underline{\sigma}$ satisfying:

$$\pi_2(z_b;\underline{\sigma}) = \pi^\star < \pi_2(z_g;\underline{\sigma}). \tag{11}$$

[10] Unless otherwise specified, in this and the remaining figures, the parameters used are: $\pi_1 = 0.7$, $v_L = 1$, $v_H = 2$, $\delta = 0.85$, and $\lambda = 0.9$.

Note that $\sigma = \underline{\sigma}$ is consistent with any price in the interval $[\underline{p},\ v_H]$, where:

$$\underline{p} \equiv v_H - \delta\rho_H(z_b;\underline{\sigma})(v_H - v_L), \tag{12}$$

and where $\rho_H(z_b;\sigma) = 1 - \lambda\sigma$ by (9). Since the seller is indifferent between offering v_H and v_L after *bad news* in the second period, he can mix between them.[11] As the probability $\phi(z_b)$ of the high offer decreases, the seller must decrease the first-period offer to keep the H-types indifferent. The lowest offer $\widetilde{p}_1 = \underline{p}$ arises when v_H is offered with probability zero in the second period following *bad news*.

Second, for any first-period offer below $v_H - \delta(v_H - v_L)$, the H-type buyers will accept the offer with probability one, since the price will never be below v_L in the second period. When the price is equal to $v_H - \delta(v_H - v_L)$, the H-type buyers are indifferent between trading and not trading, as long as they expect a price v_L in the second period with certainty. This belief is consistent with equilibrium play only if the seller's posterior is sufficiently low, i.e., $\pi_2(z;\sigma) \leq \pi^\star$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \geq \bar{\sigma}$ for $\bar{\sigma}$ satisfying:

$$\pi_2(z_b;\bar{\sigma}) < \pi^\star = \pi_2(z_g;\bar{\sigma}). \tag{13}$$

Note that $\sigma = \bar{\sigma}$ is consistent with any price in the interval $[v_H - \delta(v_H - v_L),\ \bar{p}]$, where:

$$\bar{p} \equiv v_H - \delta\rho_H(z_b;\bar{\sigma})(v_H - v_L). \tag{14}$$

Since the seller is indifferent between offering v_H and v_L following *good news* in the second period, he can again mix.[12] As the probability $\phi(z_g)$ of the high offer increases, the seller must increase the first-period offer to keep the H-types indifferent. The highest offer $\widetilde{p}_1 = \bar{p}$ arises when v_H is offered with probability one in the second period following *good news*.

Finally, when $\sigma \in (\underline{\sigma}, \bar{\sigma})$, the seller's posterior beliefs satisfy:

$$\pi_2(z_b;\sigma) < \pi^\star < \pi_2(z_g;\sigma), \tag{15}$$

so that his offers in the second period are v_H following *good news* and v_L following *bad news*. This in turn gives rise to strategic complementarity in the buyers' trading behavior. As σ rises, high-valuation buyers expect good news to arrive more frequently in the second period and are therefore willing to accept a higher offer in the first period. This leads to the upward-sloping section in the (inverse) demand schedule depicted in Figure 1. Consequently, for any offer $\widetilde{p}_1 \in [\underline{p}, \bar{p}]$, the demand schedule is a correspondence that consists of the trading probabilities $\underline{\sigma}$, $\bar{\sigma}$, and $\sigma^\star(\widetilde{p}_1)$,[13] where $\sigma^\star(\widetilde{p}_1)$ satisfies the H-type's indifference condition:

$$v_H - \widetilde{p}_1 = \delta\rho_H\left(z_b;\sigma^\star(\widetilde{p}_1)\right)(v_H - v_L). \tag{16}$$

We highlight that the nonmonotonicity of the inverse demand depicted in Figure 1 arises solely due to information spillovers. When valuations are uncorrelated, a buyer's trading behavior conveys no information

[11] Since the posterior belief following good news is strictly above $\pi^\star$, the offer in that event is v_H.

[12] Since the posterior belief following bad news is strictly below $\pi^\star$, the offer in that event is v_L.

[13] These trading probabilities are distinct for $\tilde{p}_1 \in (\underline{p}, \bar{p})$ and coincide pairwise at the endpoints ($\sigma^\star(\underline{p}) = \underline{\sigma}$ and $\sigma^\star(\bar{p}) = \bar{\sigma}$).

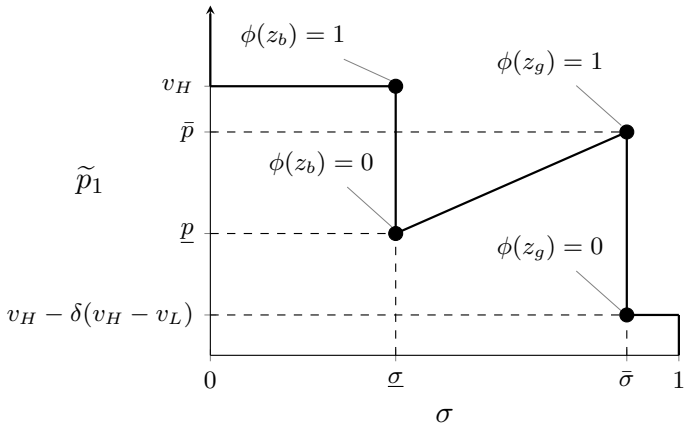


Figure 1: The inverse demand schedule with information spillovers ($\lambda > \pi_1$).

about the other buyer's valuation; thus, $\underline{\sigma} = \bar{\sigma}$ and the inverse demand is monotonic, as depicted in Figure 2. As correlation increases, the gap $\bar{\sigma} - \underline{\sigma}$, and hence the nonmonotonicity region, expands. On the one hand, the threshold $\underline{\sigma}$ decreases with λ because higher correlation makes the seller's posterior belief π_2 *more pessimistic* following *bad news*, requiring high-valuation buyers to trade *less aggressively* in the first period to keep the posterior at $\pi^\star$. On the other hand, $\bar{\sigma}$ increases in λ because higher correlation makes the belief π_2 *more optimistic* following *good news*, requiring high-valuation buyers to trade *more aggressively* in the first period to keep the posterior at $\pi^\star$.

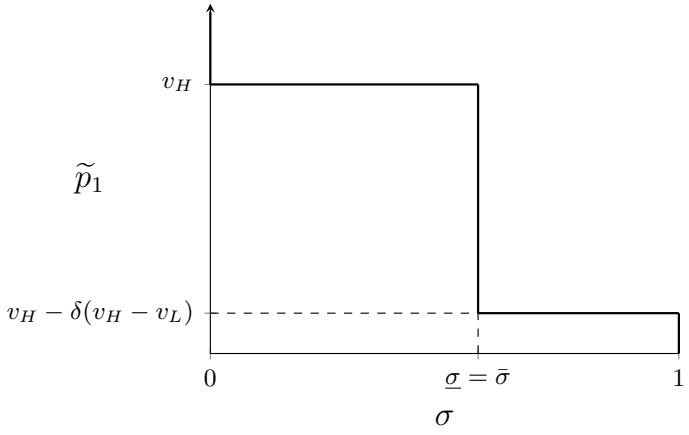


Figure 2: The inverse demand schedule absent information spillovers ($\lambda = \pi_1$).

We summarize the key features of the demand schedule discussed above in the following proposition, which we will use shortly for equilibrium characterization.

Proposition 2. *With information spillovers, i.e., $\lambda > \pi_1$, the inverse demand schedule is nonmonotone. Specifically, as the first-period price p_1 rises from $\underline{p}$ to $\bar{p}$, the trading probability $\sigma(p_1)$ takes values in $\{\underline{\sigma}, \sigma^\star(p_1), \bar{\sigma}\}$, where $\sigma^\star(p_1)$ increases from $\underline{\sigma}$ to $\bar{\sigma}$. Moreover, the region of nonmonotonicity expands with the strength of spillovers: whereas the threshold $\bar{\sigma}$ increases in λ, the threshold $\underline{\sigma}$ decreases in it.*

3.2 From Demand to Equilibria

Given the characterization of the demand schedule, for any first-period offer p_1 and an associated continuation game summarized by $\{\sigma(p_1), \phi(b, p_1), \phi(g, p_1)\}$, we can compute the seller's corresponding expected revenue. Consequently, we can characterize the set of the seller's optimal pricing strategies and the associated trading probabilities of the buyers. As a first step, we classify potential equilibria into three distinct categories.

Proposition 3. *In any equilibrium, the first-period offer and the trading probabilities fall into one of the following categories:*

- ***High price:*** $p_1 = v_H$, $\sigma(p_1) = \underline{\sigma}$, *and the L-type never trades.*
- ***Low price:*** $p_1 = v_H - \delta(v_H - v_L)$, $\sigma(p_1) = 1$, *and the L-type trades after every news.*
- ***Contingent price:*** $p_1 \in [\underline{p}, \bar{p}]$, $\sigma(p_1) \in \{\sigma^\star(p_1), \bar{\sigma}\}$ *for* $\sigma^\star(\cdot)$ *defined by* (16)*, and the L-type trades with probability one after bad news and with probability less than one after good news.*

Any first-period offer above v_H is strictly dominated by charging a price slightly below v_H and ensuring that H-types accept with positive probability. Likewise, any first-period offer in the interval $(\bar{p}, v_H)$ is dominated by a slightly higher offer $\tilde{p}_1$ within the same interval, adjusting $\phi(b, \tilde{p}_1)$ so that the first-period trading probability remains $\underline{\sigma}$. The trading probability and the second-period outcome remain unchanged, but the seller's revenue is higher under $\tilde{p}_1$. At $p_1 = v_H$, no equilibrium exists with $\sigma < \underline{\sigma}$, since the seller can profitably deviate to a slightly lower offer—thereby increasing the first-period trading probability without meaningfully reducing the price or altering second-period outcomes.

By the same reasoning, any first-period offer in the interval $\big(v_H - \delta(v_H - v_L), \underline{p}\big)$ is strictly dominated by a slightly higher offer $\tilde{p}_1$ within the same interval, adjusting $\phi(g, \tilde{p}_1)$ accordingly to keep the first-period trading probability at $\bar{\sigma}$. From Lemma 1, there is no equilibrium in which $p_1 < v_H - \delta(v_H - v_L)$. By logic analogous to that in the previous paragraph, no equilibrium exists with $\sigma < 1$ and $p_1 = v_H - \delta(v_H - v_L)$.

No equilibrium has $p_1 \in [\underline{p}, \bar{p}]$ with $\sigma(p_1) = \underline{\sigma}$: on this branch the seller is indifferent after bad news, so his per-buyer revenue is $\pi_1 \underline{\sigma} p_1 + \delta \pi_1 (1 - \underline{\sigma}) v_H$ whatever $\phi(b, p_1)$, and any offer in $(\bar{p}, v_H)$, where $\underline{\sigma}$ is the only consistent trading probability, yields strictly more.

We refer to equilibria with first-period offers in the interval $[\underline{p}, \bar{p}]$ as *contingent-price* equilibria (just as in Proposition 1), since the seller charges both v_H and v_L with positive probability following different news realizations. For every p_1 within this interval, we have $\phi(g, p_1) > 0$ and $\phi(b, p_1) < 1$ (see Figure 1). In this case, the seller uses information from first-period trades to better separate buyer types in the second period.

We are now ready to state our main result.

Theorem 1. *An equilibrium exists and, for given δ, there is a threshold $\lambda^{nc} \in (\pi_1, 1)$ such that:*

(i) If $\lambda < \lambda^{nc}$, the equilibrium is generically unique and it is either a ***high-price*** *or a* ***low-price*** *equilibrium.*

(ii) If $\lambda > \lambda^{nc}$, there is a continuum of ***contingent-price*** *equilibria, which may coexist with either a* ***high-price*** *or a* ***low-price*** *equilibrium.*

When buyers' valuations are independent, the monopolist treats each buyer as a separate screening problem. Because one buyer's decision conveys no information about the other, the seller's second-period pricing is unaffected by early trades. His problem, therefore, reduces to the familiar dynamic monopoly trade-off between extracting more surplus from high-valuation buyers and selling to low-valuation ones. The equilibrium is unique: either a high-price equilibrium, in which only high types buy gradually over time, or a low-price equilibrium, in which the seller cuts prices to induce low types to trade in the second period. This benchmark corresponds to classic models of dynamic price discrimination with private information and no cross-buyer learning.[14]

Once valuations are correlated, trading decisions reveal information across buyers. A sale in the first period signals that the remaining buyer is more likely to have a high valuation, while the absence of a sale suggests the opposite. These information spillovers alter the seller's incentives: his optimal second-period price may now depend on what has been inferred from early trades. In particular, he may set a high price in the second period only after an early sale. Buyers' strategies then exhibit strategic complementarities: each buyer's willingness to trade in the first period increases with the perceived likelihood that others will do so as well. Hence, higher prices can be sustained together with higher trade, both reinforced by beliefs that others are also trading.

As shown in Figure 1 and Proposition 2, as correlation strengthens, the seller faces larger regions where raising the price can increase the probability of trade. For low correlation, this effect is too weak to matter, and the equilibrium remains unique as in the uncorrelated benchmark. Beyond a correlation threshold (i.e., $\lambda > \lambda^{nc}$), however, the feedback between prices, beliefs, and trading behavior supports multiple self-fulfilling outcomes. The seller may post a single high or low price, or condition future prices on the information revealed by early sales, generating a continuum of contingent-price equilibria. Correlation thus transforms a standard dynamic screening problem into one in which information links buyers' strategies, shapes beliefs, and generates equilibrium multiplicity.

The equilibrium set—illustrated in Figure 3—is constructed as follows. In any equilibrium, the seller posts a first-period price p_1, which is accepted with probability $\sigma(p_1)$. For instance, the high-price equilibrium corresponds to $(p_1, \sigma(p_1)) = (v_H, \underline{\sigma})$. Suppose the seller deviates to some offer $\widetilde{p}_1 \neq p_1$. Following such a deviation, both buyers believe that the other buyer will accept the offer with the lowest acceptance probability $\sigma(\widetilde{p}_1)$ consistent with the demand schedule. This off-path belief has bite only for deviations in the interval $[\underline{p}, \bar{p}]$, where inverse demand is nonmonotone and information spillovers enable buyers to "coordinate" on a continuation that minimizes the seller's revenue. Given this structure, the equilibrium set follows directly.

The high-price equilibrium exists whenever its induced revenue exceeds what the seller would obtain by deviating either to a price just below $v_H - \delta(v_H - v_L)$—which induces certain first-period trade by high types—or to a price just below $\underline{p}$, which yields trade with probability $\bar{\sigma}$.[15] As shown in Panel B of the figure, this condition holds when the seller is sufficiently patient: a high discount factor δ raises the value of future rents and allows him to sacrifice early trade volume to extract full surplus from high types.

[14] See Sobel and Takahashi (1983), Fudenberg and Tirole (1983), and Fuchs and Skrzypacz (2013).

[15] Following any deviation to the interval $[\underline{p}, \bar{p}]$, the seller's most unfavorable continuation is one in which high types trade in the first period with probability $\underline{\sigma}$. In this continuation, the seller gets strictly less than v_H in the first period while facing the same trade probability as in the high-price equilibrium.

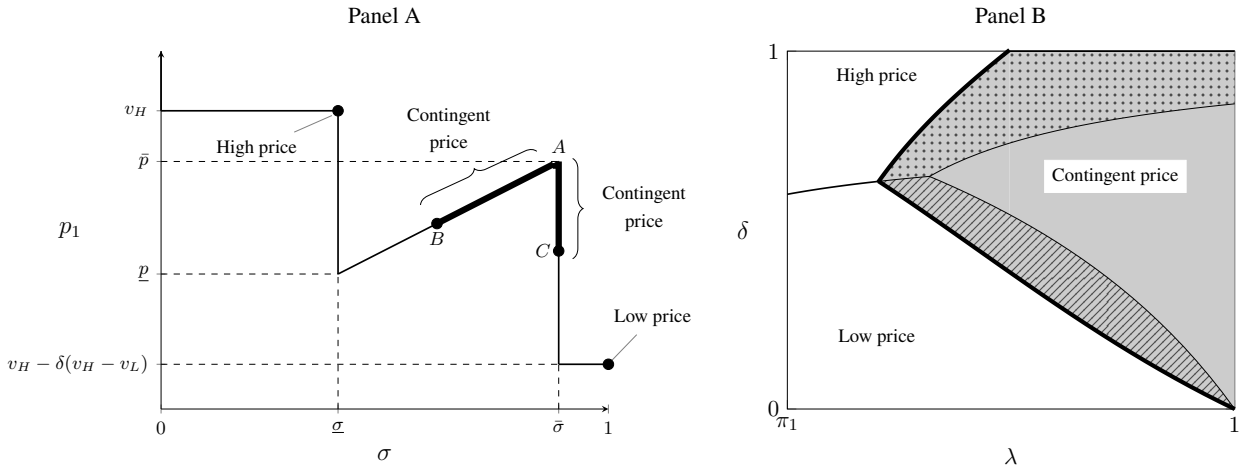


Figure 3: **Equilibrium and coexistence**. Panel A depicts equilibrium prices and trading probabilities for a given set of parameters, whereas Panel B depicts the equilibrium set as a function of λ and δ. In Panel B, the gray-shaded area is where contingent-price equilibria exist; its thick boundary depicts threshold λ^{nc} as a function of δ. In the dotted area, contingent-price equilibria coexist with the high-price equilibrium; in the striped area they coexist with the low-price equilibrium.

Similarly, the low-price equilibrium exists when its revenue exceeds that from deviating either to a price just below v_H—which yields trade with probability $\underline{\sigma}$—or to a price just below $\underline{p}$, which yields trade with probability $\bar{\sigma}$. As shown in Panel B, this occurs when the seller is relatively impatient: a low δ raises the value of early sales, inducing the seller to prioritize quick trade over surplus extraction.

Finally, a contingent-price equilibrium exists whenever its revenue exceeds that from deviating either to a price just below $v_H - \delta(v_H - v_L)$, which ensures trade with probability one, to a price just below $\underline{p}$, which yields trade with probability $\bar{\sigma}$, or to a price just below v_H, which yields trade with probability $\underline{\sigma}$. Such equilibria arise when information spillovers are strong enough to relax the seller's trade-off between charging high prices and sustaining first-period trade. While strategic complementarities among buyers can support equilibria with high prices (e.g., point A in Panel A), they also sustain equilibria in which buyers anticipate low trade following a deviation to high prices, inducing the seller to set low prices (e.g., point C).

3.3 Efficient Equilibria

Thus far, we have shown that—absent seller commitment—strong information spillovers give rise to equilibrium multiplicity. As we will see shortly, this multiplicity has important implications for how information spillovers affect social surplus and its allocation between the seller and the buyers. Before turning to these implications, we also characterize the set of Pareto-efficient equilibria, defined as those for which no other equilibrium makes both the buyers and the seller weakly better off, with at least one group strictly so. Examining this set provides a useful way to organize equilibrium outcomes and to highlight whether information spillovers give rise to a conflict between efficiency and surplus extraction.

Proposition 4. *If $\lambda > \lambda^{nc}$, the Pareto-efficient equilibrium set consists of essentially all contingent-price equilibria with $\sigma = \bar{\sigma}$, and the low-price equilibrium if the latter exists.* [16]

[16] We say essentially because the contingent-price equilibrium with $\sigma = \bar{\sigma}$ and the lowest first-period offer, in which the seller

When an equilibrium is unique, the Pareto-efficient set trivially includes it. When multiple equilibria coexist, the ranking depends on the nature of coexistence. First, consider the case in which the high-price equilibrium coexists with a contingent-price equilibrium (the dotted-shaded region in Panel B of Figure 3). The latter must Pareto dominate the former. Buyers are better off: they obtain a positive surplus in the contingent-price equilibrium but zero surplus in the high-price equilibrium. The seller is also weakly better off. By revealed preference—since the contingent-price equilibrium exists—he prefers to accelerate trade with H-types and to trade with L-types following bad news, rather than post $v_H - \varepsilon$ and obtain essentially the same payoff as in the high-price equilibrium for small ε.

Second, consider the case in which multiple contingent-price equilibria coexist (the shaded region in Panel B). Here, any contingent-price equilibrium with first-period offer p_1 and trading probability $\sigma = \sigma^\star(p_1)$—where $\sigma^\star(\cdot)$ lies on the upward-sloping segment of the demand schedule (see Panel A)—is dominated by another contingent-price equilibrium with the same offer but a higher trading probability $\sigma = \bar{\sigma}$. Buyers earn the same surplus in both equilibria, while the seller is strictly better off as he accelerates trade with H-types.

Finally, although contingent-price equilibria with $\sigma = \bar{\sigma}$ dominate those with $\sigma = \sigma^\star$, no Pareto ranking exists among contingent-price equilibria with different first-period offers but the same $\sigma = \bar{\sigma}$. Higher first-period prices—reflecting a greater likelihood of charging v_H after good news—increase the seller's revenue but reduce the surplus of H-type buyers. Likewise, there is no Pareto ordering between contingent-price equilibria and low-price equilibrium: buyer surplus is maximal in the low-price equilibrium, whereas the seller's revenue is weakly higher in all contingent-price equilibria than in the low-price equilibrium.

4 Information, Commitment, and Welfare

We now build on our equilibrium characterization to study how information spillovers affect agents' payoffs. Because spillovers can generate equilibrium multiplicity (Theorem 1), a given parameter configuration may be associated with multiple trading outcomes and payoff profiles. We therefore characterize the effect of spillovers over the entire equilibrium set. In particular, we study how information spillovers affect the seller's ability to extract surplus from buyers—through his effective commitment power to sustain high prices—and how they shape social surplus, defined as the present discounted gains from trade. A central theme of this section is that stronger spillovers can either undermine or foster the seller's commitment to maintain high prices; moreover, once multiple equilibria emerge, the seller's and the buyers' payoffs may either be aligned or move in opposite directions as spillovers intensify.

4.1 Spillovers and Seller Commitment

We begin by studying how information spillovers affect the seller's equilibrium payoff, defined as his present discounted expected revenue. Given an equilibrium characterized by a first-period price-quantity pair (p_1, σ), the seller's per-buyer expected revenue is:

earns exactly his low-price revenue, is Pareto dominated by the low-price equilibrium when the latter exists.

$$R(p_1, \sigma; \lambda) = \pi_1 \left(\sigma p_1 + (1-\sigma)\delta \sum_{z \in \{z_b, z_g\}} \rho_H(z; \sigma) \left[\phi(z) v_H + (1 - \phi(z)) v_L \right] \right) + (1 - \pi_1)\delta \sum_{z \in \{z_b, z_g\}} \rho_L(z; \sigma)(1 - \phi(z)) v_L, \tag{17}$$

where the news distribution $\{\rho_\theta(z;\sigma)\}$ is given in (9). We note that the pair (p_1, σ) uniquely determines the probabilities $\phi(z_g)$ and $\phi(z_b)$ with which the seller makes a high second-period offer following good and bad news, respectively (see Figure 1).

The first term in (17) captures the surplus extracted from high-valuation buyers, while the second term captures the surplus extracted from low-valuation buyers. Spillovers affect the seller's revenue both directly—by changing the informativeness of news through the correlation parameter λ—and indirectly, through their impact on first-period trade σ and price p_1, and on the seller's ability to condition continuation prices on observed news.

Given a correlation λ, let $\mathcal{R}(\lambda)$ denote the set of revenue levels attained across equilibria, and let $\overline{R}(\lambda)$ and $\underline{R}(\lambda)$ denote its maximal and minimal elements, respectively.

Proposition 5. *The correspondence $\lambda \mapsto \mathcal{R}(\lambda)$ is hemi-continuous. Moreover:*

1) For $\lambda < \lambda^{nc}$, the unique revenue level $\underline{R}(\lambda) = \overline{R}(\lambda)$ is weakly decreasing in λ.

2) For $\lambda \geq \lambda^{nc}$:

(i) The lowest attainable revenue, $\underline{R}(\lambda)$, can be nonmonotone in λ and remains bounded away from the seller's payoff under commitment for all λ.

(ii) The highest attainable revenue, $\overline{R}(\lambda)$, is strictly increasing in λ and converges to the seller's payoff under commitment as λ goes to 1.

Figure 4 illustrates Proposition 5 and contrasts it with the commitment benchmark. Under commitment to second-period offers—our benchmark in Section 2.2—spillovers benefit the seller, as shown by the x-marked locus in the figure. As Proposition 1 establishes, when $\lambda < \lambda^c$ the committed seller's payoff is invariant to spillovers, since he optimally disregards the news when setting prices. Once λ exceeds λ^c, the seller conditions his second-period offer on the news realization, thereby imposing less favorable continuation terms on the high type. As correlation increases, this continuation threat becomes more effective at disciplining delay, enabling the seller to extract a larger share of surplus in the first period. In the limit as $\lambda \to 1$, the seller extracts the full surplus, albeit with delay.

Absent commitment, the effects of information spillovers change markedly, and it is useful to distinguish between two cases, depending on whether the equilibrium is unique. Suppose first that spillovers are weak, so that $\lambda < \lambda^{nc}$ and the equilibrium is unique. Consider the parameter region in which equilibrium play features

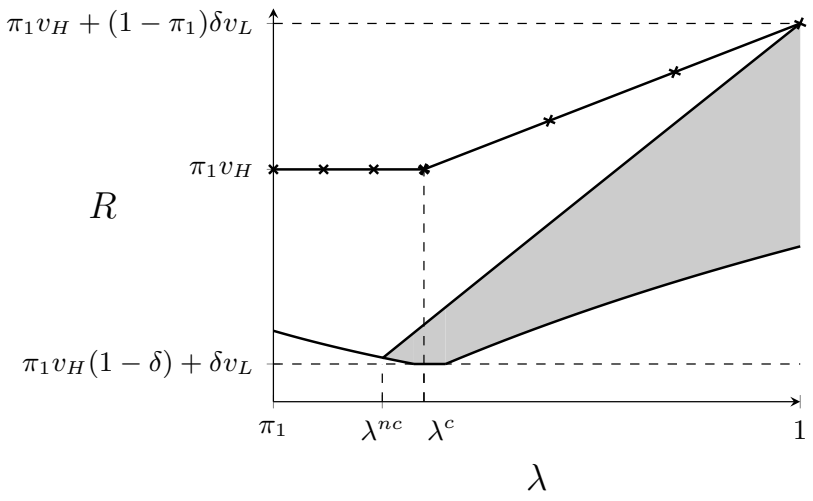


Figure 4: **Seller's revenue**. The figure depicts the seller's revenue with commitment (x-marked locus) and across all equilibria without commitment (solid line for $\lambda < \lambda^{nc}$ and shaded region between solid loci otherwise). Here, we set $\delta = 0.645$. For this parameterization, contingent pricing arises without commitment for $\lambda \in (\lambda^{nc}, \lambda^c)$, where a committed seller would ignore the news.

high prices, with $p_1 = v_H$, $\sigma = \underline{\sigma}$, and $\phi(z_g) = \phi(z_b) = 1$ (Figure 4).[17] In this region, revenue is decreasing in λ. The reason is that stronger spillovers make the seller's posterior after bad news more pessimistic. In the absence of commitment, sustaining high second-period prices then requires reducing first-period trade.[18] Thus, stronger spillovers *undermine* the seller's effective commitment power: the more informative future signals are, the more the seller must distort current trade to sustain high future prices.[19]

If correlation increases beyond λ^{nc}, however, a new class of equilibria with contingent pricing emerges (Theorem 1). In these equilibria, the seller conditions second-period prices on realized news and can sustain high prices following good signals without sharply restricting first-period trade. Outcomes that were previously infeasible become supportable in equilibrium. In this sense, stronger spillovers can also *foster* effective commitment.

Now suppose $\lambda > \lambda^{nc}$, so that there are multiple equilibria (Theorem 1). In this region, comparative statics are equilibrium-dependent. Along the revenue-maximizing equilibrium, the seller sets contingent second-period prices. Because he already charges a low price following bad news, a further increase in λ does not erode his ability to sustain a high price after good news. Moreover, as λ increases, a high-type buyer assigns a higher probability to the other buyer also being a high type and therefore anticipates a greater likelihood of facing a high second-period price. This relaxes the first-period incentive constraint: high types are willing to accept a higher initial price, allowing the seller to raise p_1 and increase revenue (as shown by the upper boundary of the shaded region in Figure 4). As $\lambda \to 1$, revenue in this equilibrium monotonically converges to the commitment payoff.[20]

[17] A necessary and sufficient condition is that revenue from posting $p_1 = v_H$, equal to $\pi_1 v_H[\underline{\sigma} + (1-\underline{\sigma})\delta]$, exceeds revenue from posting $p_1 = v_H - \delta(v_H - v_L)$, equal to $\pi_1[v_H - \delta(v_H - v_L)] + (1-\pi_1)\delta v_L$. This holds, for example, when δ is relatively high, as depicted in Panel B of Figure 3.

[18] In this equilibrium, the first-period trading probability $\underline{\sigma}$ decreases with λ (Proposition 2).

[19] By contrast, in the other parameter region in which the unique equilibrium features low prices, i.e., $p_1 = v_H - \delta(v_H - v_L)$, $\sigma = 1$, and $\phi(z_g) = \phi(z_b) = 0$, revenue is invariant to λ.

[20] The revenue-maximizing equilibrium features $p_1 = v_H - \delta\rho_H(z_b; \bar{\sigma})(v_H - v_L)$, $\sigma = \bar{\sigma}$, $\phi(z_g) = 1$, and $\phi(z_b) = 0$, where $\bar{\sigma}$

By contrast, along revenue-minimizing equilibria, spillovers may further *undermine* effective commitment by making buyer coordination against high prices more effective. Whereas, in the revenue-maximizing equilibrium, spillovers discipline buyers, here they instead allow buyers to coordinate on low trade following a deviation to a higher price. Anticipating this response, the seller either keeps the first-period price low or raises it and accepts low trade.[21] Along these equilibria, the effect of λ on seller revenue is nonmonotone. Revenue decreases with λ when the high-price equilibrium exists (low λ in Figure 4), is constant in λ when the low-price equilibrium exists (intermediate λ), and may either rise or fall with λ when neither the high- nor the low-price equilibrium is sustainable (high λ).[22]

Thus, information spillovers expand the set of feasible outcomes and, depending on equilibrium play, can either foster or undermine the seller's effective commitment power. When spillovers are weak ($\lambda < \lambda^{nc}$), they are weakly detrimental to the seller. When spillovers are strong ($\lambda > \lambda^{nc}$), they enlarge the equilibrium set to include outcomes close to the commitment benchmark; whether this potential is realized, however, depends on equilibrium play. Importantly, these conclusions do not rely on Pareto-dominated equilibria: the revenue-maximizing equilibrium is Pareto efficient; one revenue-minimizing equilibrium maximizes buyer surplus, and intermediate revenue levels are also within the Pareto-efficient set.

4.2 Spillovers and the Distribution of Surplus

We now turn to the effects of information spillovers on social surplus and its allocation. Since in equilibrium a low-type buyer earns zero surplus, and a high-type buyer must be indifferent between accepting the first-period offer p_1 and waiting, per-buyer social surplus can be written as:

$$W(p_1, \sigma; \lambda) = R(p_1, \sigma; \lambda) + \pi_1(v_H - p_1), \tag{18}$$

where the second term is the expected rent of high-type buyers. Thus, social surplus equals the seller's expected revenue plus the information rent left to buyers.

Given correlation λ, let $\mathcal{W}(\lambda)$ be the set of welfare levels attained across equilibria, and let $\overline{W}(\lambda)$ and $\underline{W}(\lambda)$ be its maximal and minimal elements, respectively.

Proposition 6. *The correspondence $\lambda \mapsto \mathcal{W}(\lambda)$ is not lower hemi-continuous. Moreover:*

(i) The lowest attainable surplus, $\underline{W}(\lambda)$, remains bounded away from the social surplus under commitment as λ goes to 1.

(ii) The highest attainable surplus, $\overline{W}(\lambda)$, converges to the social surplus under commitment as λ goes to 1.

(iii) The social surplus along any sequence of Pareto-efficient equilibria converges to the social surplus under commitment as λ goes to 1.

---

increases with λ (Proposition 2).

[21] There are up to three revenue-minimizing equilibria. Two are contingent-price equilibria (points B and C in Panel A of Figure 3). By Theorem 1, these may coexist either with the high-price equilibrium (for low λ in Figure 4) or with the low-price equilibrium (for intermediate λ in the figure).

[22] It rises in our parameterization, but falls, e.g., for $\pi_1 = 0.98$, $v_L = 0.8$, $v_H = 2$ and $\delta = 0.5$.

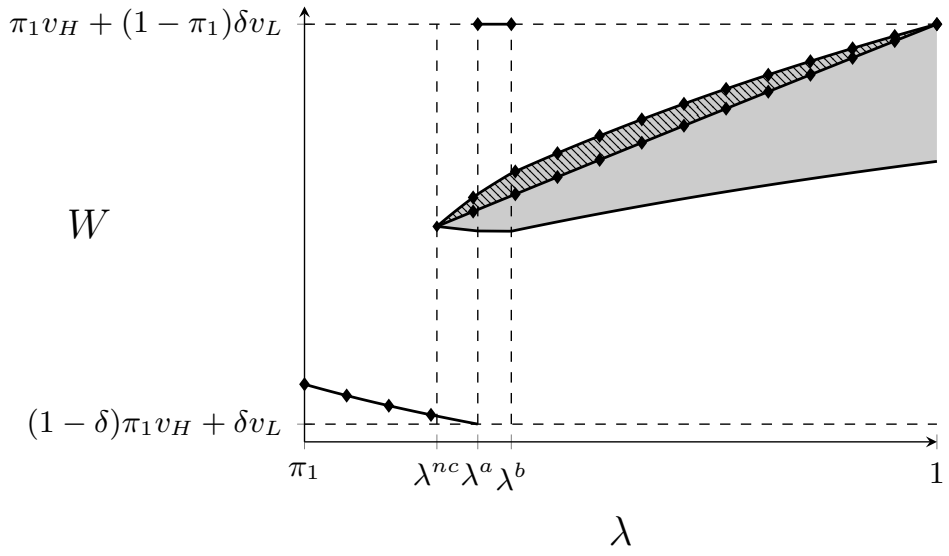


Figure 5: **Social surplus**. The figure depicts the social surplus across all equilibria without commitment (solid line for $\lambda < \lambda^{nc}$ and shaded region between solid loci otherwise). The figure uses the same parameters as Figure 4.

Figure 5 illustrates these results and contrasts them with the revenue set characterized in Proposition 5. The key difference is that the welfare set $\mathcal{W}(\lambda)$ varies discontinuously with λ, whereas the revenue set evolves continuously. These discontinuities arise at values of λ where the equilibrium set changes—either because a new class of equilibria emerges or because an existing one disappears. By contrast, the revenue set $\mathcal{R}(\lambda)$ evolves continuously: because the seller chooses p_1, equilibrium switching requires him to be indifferent across continuation outcomes, which smooths revenue at transition points. No analogous force disciplines welfare, as different equilibria involve discrete changes in p_1, and thus in buyer rents (Equation (18)).

For $\lambda < \lambda^{nc}$, under the parameterization of Figure 5, the unique equilibrium features a high price, $p_1 = v_H$. In this case, social surplus coincides with the seller's revenue (Equation (18)), so the effect of information spillovers on welfare mirrors their effect on revenue.

As λ crosses the threshold λ^{nc}, contingent-price equilibria emerge. At this point, the seller is indifferent between the high price and an intermediate price that induces news-dependent pricing in the continuation game. In all contingent-price equilibria, however, the first-period price is strictly lower than in the high-price equilibrium, generating strictly positive buyer surplus. Since the high-price equilibrium continues to exist, the welfare set expands to include strictly higher values; that is, $\overline{W}(\lambda)$ jumps above $\underline{W}(\lambda)$ at $\lambda = \lambda^{nc}$.

As λ increases further and crosses λ^a, the correlation at which the high-price and the low-price revenues coincide, the high-price equilibrium disappears and the low-price equilibrium emerges. Although the seller remains indifferent across equilibria at the switching point, the low-price equilibrium delivers strictly higher buyer surplus than all other equilibria. Hence, the welfare frontier shifts upward, and $\overline{W}(\lambda)$ jumps at $\lambda = \lambda^a$.

Finally, as λ rises further and crosses λ^b, the correlation at which the low-price revenue equals the lowest contingent-price revenue, the low-price equilibrium also disappears, leaving only contingent-price equilibria. At this point, the welfare set contracts discontinuously, with $\overline{W}(\lambda)$ jumping downward at $\lambda = \lambda^b$. For all higher values of λ, only contingent-price equilibria exist and the welfare set evolves continuously with λ.

Therefore, as with seller revenue, the set of attainable welfare outcomes expands when information spillovers become sufficiently strong. Along the welfare-maximizing equilibrium—corresponding either to the low-price equilibrium or to the contingent-price equilibrium with the lowest first-period price (point C in Panel A of Figure 3)—social surplus converges to its commitment benchmark as $\lambda \to 1$. By contrast, along the welfare-minimizing equilibrium—corresponding, in the figure's parameterization, either to the high-price equilibrium or to the contingent-price equilibrium with intermediate trade $\sigma^\star$ (point B in Panel A of Figure 3)—social surplus remains bounded away from its commitment level. However, if attention is restricted to Pareto-efficient equilibria (the striped region in Figure 5), even the lowest welfare outcomes converge to their commitment counterparts as λ approaches one.

The preceding analysis shows that information spillovers reshape both the level and the distribution of surplus across equilibria. Although spillovers affect revenue and welfare through the same pricing decisions, revenue and welfare need not move in tandem: changes in the first-period price p_1 reallocate surplus between the seller and buyers, even when the seller remains indifferent. Consequently, spillovers may either raise or lower social surplus relative to revenue, depending on how they affect the equilibrium set.

To explore this further, consider the following perspective. Taking as a benchmark the case without spillovers, i.e., $\lambda = \pi_1$, we say that, in a given equilibrium, spillovers λ benefit (harm) the seller if revenue $R(p_1, \sigma; \lambda)$ is higher (lower) than in the benchmark.[23] Analogous definitions apply to buyers and social surplus. This perspective is particularly useful in applications discussed in Section 5, where spillovers arise endogenously from firms' actions (e.g., investments in data). Our first result identifies when spillovers generate a conflict between the seller, on the one hand, and the buyers and social efficiency, on the other.

Proposition 7 (Conflict)**.** *Suppose that the low-price equilibrium is played absent information spillovers, i.e., when $\lambda = \pi_1$. If information spillovers benefit the seller, then they harm both the buyers and social surplus.*

This result follows from the fact that the low-price equilibrium maximizes both buyer surplus and social surplus among all equilibria for any λ. For spillovers to benefit the seller, it must be that they induce a shift toward contingent-price equilibria (which occurs for $\lambda > \lambda^{nc}$), involving higher first-period prices and reduced trade. These changes strictly lower both buyer surplus and social surplus. However, as our next result shows, information spillovers may also generate alignment between the seller and the buyers.

Proposition 8 (Alignment)**.** *Suppose that the high-price equilibrium is played absent information spillovers, i.e., when $\lambda = \pi_1$. If information spillovers benefit the seller, then they also benefit the buyers and social surplus.*

This result is also intuitive. In the high-price equilibrium, buyers receive zero surplus. Moreover, its revenue decreases with λ because $\underline{\sigma}$ does (Proposition 2). Thus, in that equilibrium, the seller captures the entire surplus, and that surplus is locally decreasing in λ, as illustrated in Figure 5. For spillovers to benefit the seller, they must induce a transition to a contingent-price equilibrium, in which buyers obtain strictly positive surplus. In that case, which can only occur for $\lambda > \lambda^{nc}$, buyer surplus and social surplus rise together with seller revenue.

[23] As multiple equilibria exist when $\lambda > \lambda^{nc}$, it is possible that spillovers benefit the seller in one equilibrium but not in another.

As we discuss next, these results provide insights into why firms may want to invest in data, experiment with consumers, or facilitate coordination among their input suppliers. At the same time, they highlight that increased access to consumer data need not be socially desirable, as it may strengthen information spillovers and shift the equilibrium set toward outcomes that favor surplus extraction over efficiency.

5 Applications

We next apply our framework to three economically important settings: data investment, A/B testing, and supplier coordination. In each application, firms strategically shape the information generated by trade, thereby affecting future pricing incentives, market outcomes, and welfare.

5.1 Data Investments

Recent advances in data processing have substantially expanded firms' use of consumer data. Survey evidence indicates that the share of U.S. medium-to-large firms systematically incorporating data into business decisions has risen sharply over the past two decades (Brynjolfsson and McElheran 2016). A major concern is that richer consumer data enables increasingly refined forms of price discrimination, a concern supported by evidence of systematic variation in online prices and targeted promotions across consumer groups (Hannak et al. 2014; Adams et al. 2025).

A defining feature of these practices is that firms use consumer information to segment buyers into increasingly homogeneous groups. In our framework, this corresponds to increasing the correlation of valuations within groups and thereby strengthening the information spillovers generated by trade. To capture this channel, we allow the seller to invest in data at cost $c > 0$ at $t = 0$, which raises the correlation parameter from π_1 to $\lambda^D > \pi_1$, with the investment publicly observed before trading begins at $t = 1$.

Corollary 1. *There exists a threshold $\bar{\lambda}^D$ such that the seller invests in data only if $\lambda^D > \bar{\lambda}^D$. If investment occurs, buyer surplus and social surplus may either increase or decrease.*

The intuition follows directly from our earlier characterization. Consider first the case in which, absent spillovers, the market features the low-price equilibrium. Weak increases in correlation (that is, $\lambda^D < \lambda^{nc}$) leave equilibrium pricing unchanged and therefore do not justify data investment. Strong increases in correlation (that is, $\lambda^D > \lambda^{nc}$) become profitable only to the extent that they sustain contingent-price equilibria. By Proposition 7, however, these equilibria necessarily reduce both buyer surplus and efficiency. In this case, data investment is privately profitable but socially harmful.

Consider next the case in which, absent spillovers, the market features the high-price equilibrium. When spillovers are sufficiently strong, contingent-price equilibria become sustainable and they feature more first-period trade by high types and second-period trade with low types. By Proposition 8, any data investment that benefits the seller must also increase buyer and social surplus. In this case, investment in consumer data is socially beneficial.

The policy debate surrounding consumer data has largely focused on the adverse effects of data-driven price discrimination on consumer surplus. Our results instead point to a broader trade-off. Data investment

and the resulting increase in correlation can either raise or reduce buyer surplus and *total* social surplus, depending on the magnitude of equilibrium distortions absent spillovers. When distortions are large—so that the unique equilibrium features high prices—data can improve outcomes by expanding the set of equilibria with higher trade. When distortions are small—so that the unique equilibrium features low prices—the same forces may instead intensify surplus extraction and reduce efficiency.

This finding lends support to policies, e.g., the General Data Protection Regulation (GDPR) or the Digital Services Act (DSA) in the EU, which restrict firms' ability to collect, store, and exploit consumer data.[24] At the same time, it cautions against blanket restrictions, as limiting data use is not always socially beneficial.

5.2 A/B Testing

Firms in digital markets routinely rely on online randomized experimentation—or A/B testing—to guide pricing, advertising, and platform-design decisions (Kohavi, Tang and Xu 2020).

More generally, experimentation allows firms to use a subset of consumers as informational probes for market-wide demand conditions. In our model, this corresponds to relaxing anonymity: the seller induces one buyer (buyer A) to trade early and uses the information revealed by her purchase decision to refine subsequent offers to another buyer (buyer B). A full analysis of experimentation with personalized, privately observed offers would introduce higher-order belief dynamics and substantially complicate the equilibrium analysis. Instead, we consider a simple experimentation policy that isolates the informational force.

Specifically, the seller offers buyer A a sufficiently low first-period price to induce early trade, while offering buyer B a sufficiently high price to induce delay. Buyer A thus serves as an informational probe: the seller sacrifices revenue on that transaction in exchange for information about buyer B's valuation, and conditions second-period pricing for B on A's trading behavior. Since every equilibrium without experimentation yields at least the lowest total equilibrium revenue $2\underline{R}(\lambda)$, experimentation can be profitable only if its revenue exceeds $2\underline{R}(\lambda)$. Whether it does so depends on the value of the information generated by buyer A's purchase decision, which depends on the correlation level.

Corollary 2. *There exists a threshold $\lambda^{A/B}$ such that the seller experiments only if $\lambda > \lambda^{A/B}$. If experimentation occurs, buyer surplus and social surplus may either increase or decrease.*

When correlation is low, the information is too weak to affect the seller's optimal continuation price, so experimentation cannot recoup the revenue sacrificed on buyer A. Only once correlation exceeds λ^c, the threshold of Proposition 1, does the seller optimally condition second-period prices for buyer B on buyer A's purchase history. This is not sufficient for experimentation to be profitable: the revenue from experimentation must also exceed $2\underline{R}(\lambda)$. The threshold $\lambda^{A/B} > \lambda^c$ is the correlation above which this first happens, with $\lambda^{A/B} = 1$ if it never does.

Figure 6 illustrates this threshold property by comparing the seller's minimum equilibrium revenue without experimentation (solid lines) to the revenue under the experimentation policy (x-marked lines). The two panels differ only in the equilibrium that yields the seller's minimum revenue absent experimentation.

[24] See https://gdpr.eu and https://digital-strategy.ec.europa.eu/en/policies/digital-services-act for details of the GDPR and the DSA, respectively.

In Panel A, this benchmark is the high-price equilibrium, whereas in Panel B it is initially the low-price equilibrium and, for higher values of λ, the lowest-revenue contingent-price equilibrium. In both cases, experimentation can be profitable only once the correlation is sufficiently high.

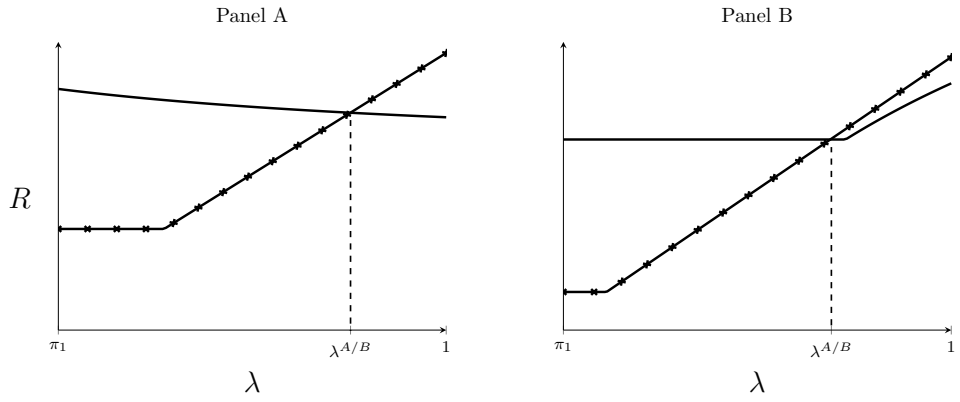


Figure 6: **Revenue with and without experimentation**. The figure depicts the seller's total revenue without A/B testing (solid lines) and with A/B testing (x-marked lines). In both panels, $v_L = 0.8$ and $v_H = 2$. In addition, in Panel A, we set $\pi_1 = 0.55$ and $\delta = 0.88$, whereas in Panel B, $\pi_1 = 0.45$ and $\delta = 0.7$.

The welfare implications follow the same logic as that of Propositions 7 and 8. When the benchmark equilibrium without experimentation is inefficient because it forgoes mutually beneficial trade, experimentation can both raise the seller's revenue and improve buyer and social surplus by inducing a more efficient allocation. By contrast, when the benchmark equilibrium already realizes all gains from trade, experimentation increases the seller's ability to price discriminate but does so by delaying trade and excluding some low-valuation buyers. Consequently, it destroys gains from trade, so buyer surplus and social surplus may decline even though seller revenue rises.

Figure 6 illustrates these two cases. In Panel A, the benchmark is the high-price equilibrium, which delays trade and excludes low-valuation buyers. Experimentation relaxes these distortions while increasing the seller's revenue, so both buyer surplus and social surplus increase. In Panel B, the benchmark is initially the low-price equilibrium, which already realizes all gains from trade. Experimentation instead delays buyer B's purchase and may exclude low-valuation buyers following good news from buyer A, reducing both buyer and social surplus.

While large firms have documented substantial revenue gains from experimentation, small firms rarely engage in A/B testing despite the decline in experimentation costs.[25] Koning, Hasan and Chatterji (2022) document that, among nearly 36,000 startups in their sample, fewer than 20% used A/B testing tools. One explanation is that small firms lack the traffic needed to generate informative experiments. In the language of our model, these firms operate in environments in which information spillovers are too weak for experimentation to pay.[26]

[25] For example, A/B testing at Microsoft's search engine generated annual increases in search-advertising revenue exceeding $100 million without sacrificing user experience (Kohavi and Thomke 2017).

[26] Although the model features only two buyers, the same force extends to an environment with many buyers (Section 6). If larger firms are interpreted as serving larger customer bases, each experimental observation becomes informative about a larger pool of remaining consumers, increasing the value of experimentation. Thus, the economic force underlying our two-buyer model suggests greater experimentation by larger firms.

5.3 Coordinating Suppliers

Suppliers often invest in cost reduction or quality improvement in ways that are especially valuable within a particular buyer–supplier relationship. When formal contracts are incomplete, these investments expose suppliers to *ex post* surplus extraction by the buyer, generating the classic hold-up problem. Apart from vertical integration, empirical evidence suggests that relational governance in repeated buyer–supplier relationships plays an important role in mitigating these investment distortions (Joskow 1987; Calzolari et al. 2021; Macchiavello and Morjaria 2023). A prominent example comes from Japanese automotive supply chains. Toyota's *Kyoryoku Kai*—supplier associations such as the *Kyohokai*—facilitate coordination and repeated interaction among suppliers and the assembler, helping sustain relationship-specific investments without full vertical integration (Dyer and Nobeoka 2000).

To relate our framework to this setting, reinterpret the seller in our model as the downstream buyer and the buyers as upstream suppliers. Relationship-specific investments by suppliers increase the value of trade to the downstream buyer but also expose suppliers to *ex post* surplus extraction. When information spillovers are sufficiently strong, our model admits multiple equilibria that differ in the extent of such surplus extraction. Some equilibria leave suppliers with little surplus, discouraging investment through the classic hold-up problem. Others leave suppliers greater rents by limiting *ex post* surplus extraction. If suppliers can coordinate on these equilibria, efficient investment incentives are restored. Because greater supplier investment also benefits the downstream buyer, the buyer may itself have incentives to facilitate such coordination. This mechanism stands in sharp contrast to the relational contracting literature (Levin 2003), where hold-up is mitigated through repeated interactions that discipline opportunistic behavior. Here, by contrast, no such long-term relationship is required: information spillovers generate a continuum of equilibria, and coordination on sufficiently supplier-favorable ones is enough to sustain efficient investment, even in a two-period model.

To capture this idea, suppose that before bargaining begins, and before learning her valuation, each supplier can make an observable relationship-specific investment at a cost $F > 0$ that scales her valuation by a factor $\alpha > 1$. Thus, an investing supplier of type $\theta \in \{H, L\}$ has valuation αv_θ. Because the investment scales valuations proportionally, the equilibrium characterization follows directly from the baseline model.

Corollary 3. *Suppose suppliers can make an observable ex ante investment that increases their valuation by a factor $\alpha > 1$ at a fixed cost $F > 0$.*

(i) If $\lambda < \lambda^{nc}$ and δ is sufficiently high, there is no symmetric equilibrium in which both suppliers invest.

(ii) If $\lambda > \lambda^{nc}$ and α is sufficiently high, there exists a symmetric equilibrium in which both suppliers invest.

Consider first the case in which information spillovers are weak ($\lambda < \lambda^{nc}$). When δ is high, the downstream buyer cannot credibly commit to leaving suppliers sufficient continuation surplus to recover the investment cost. Suppliers therefore anticipate *ex post* surplus extraction and choose not to invest. The model thus reproduces the classic hold-up problem.

Instead, when information spillovers are sufficiently strong ($\lambda > \lambda^{nc}$), multiple equilibria become sustainable. Consider the buyer-favorable contingent-price equilibrium (that is, the equilibrium favorable to

suppliers under the present interpretation). In it, suppliers retain a sufficiently large share of the gains from trade that relationship-specific investment becomes privately profitable whenever α exceeds a threshold.

Crucially, the downstream buyer also benefits from coordinating suppliers on this equilibrium, relative to the same equilibrium without investment. By encouraging *ex ante* investment, the downstream buyer raises suppliers' valuations and thus the social surplus available for future extraction (some of which he does extract). More broadly, our analysis suggests that organizational arrangements such as supplier associations can mitigate hold-up not by reallocating ownership, but by helping suppliers coordinate on equilibria that preserve relationship-specific investment incentives.

6 Extensions and Robustness

Our baseline model has two buyers, two trading periods, and binary private values. The mechanism behind our results is that current trade shapes the seller's continuation pricing through what it reveals about remaining buyers. In Appendix B (Supplemental Appendix), we argue that our mechanism continues to operate when we relax each of these simplifications, as well as Assumption 1.

Continuous valuations, many buyers, common values, and low priors. First, suppose valuations are continuously distributed and correlated through an aggregate state (Appendix B.1). Buyers then use pure strategies in equilibrium: for each first-period price, buyers above a cutoff valuation trade immediately. As in the baseline, when spillovers are strong, more expected first-period trade raises the continuation price and lowers this cutoff. Strategic complementarity, the upward-sloping region of inverse demand, and equilibrium multiplicity therefore all arise in this setting as well, so none of them is an artifact of binary types or mixed strategies. Second, with an arbitrary number of buyers (Appendix B.2), the seller conditions continuation prices on how many buyers have purchased, and each buyer's incentive to wait depends on how many others she expects to trade. As in the baseline, strong spillovers generate strategic complementarities, nonmonotone demand, and equilibrium multiplicity. Third, our results are robust to a small common-value component (Appendix B.3). When this component is large, a first-period rejection also lowers the other buyer's own expected valuation, adding a social-learning channel that we leave for future work. Finally, we show that Assumption 1 is not needed (Appendix B.4). In the low-prior case, as in the baseline, strategic complementarities and equilibrium multiplicity emerge once spillovers are sufficiently strong; contingent-price equilibria then coexist with an equilibrium in which all buyers trade immediately at v_L.

Infinite horizon. Absent spillovers, an infinite horizon gives the classic dynamic monopoly problem, the setting in which the seller's lack of commitment is most severe. High types wait for price cuts, and as $\delta \to 1$ prices converge to v_L (Fudenberg, Levine and Tirole 1985; Gul, Sonnenschein and Wilson 1986). Even here, however, strong spillovers can restore the seller's commitment power (Appendix B.5). For any $\delta < 1$, as $\lambda \to 1$ there exist equilibria in which high types trade in the first period at prices converging to v_H and low types trade at v_L in the next. The seller's revenue converges to its commitment counterpart (Proposition B.1).

The argument relies on the same complementarity as in the baseline. When high types trade aggressively, an observed purchase is a strong signal that the remaining buyer is also a high type. The seller then faces a one-buyer bargaining problem with a posterior near one and extracts almost all of her surplus. After no

trade the seller is pessimistic and cuts the price to v_L. When correlation is high, however, a high type who delays almost surely ends up in the former continuation, with negligible surplus from waiting. This sustains a first-period price close to v_H, which in turn sustains the aggressive early trade that is so informative.[27]

7 Concluding Remarks

We study dynamic monopoly pricing when buyers' valuations are correlated, so that one buyer's purchase reveals information about the valuations of others. We show that these information spillovers fundamentally reshape the seller's dynamic pricing problem by making buyers' purchase decisions strategic complements. As a result, demand becomes nonmonotone, equilibrium multiplicity emerges, and information spillovers can either strengthen or weaken the seller's effective commitment power.

Our analysis highlights that the information generated through trade is not merely a by-product of market interactions but a source of market power when sellers cannot commit. By endogenizing the information revealed by trade, we show that information spillovers reshape dynamic monopoly pricing, with implications for prices, trade, social welfare, and the allocation of surplus. Our findings suggest that information generated through market interactions plays an important role in shaping market outcomes, even when firms possess no informational advantage over consumers.

Appendix A: Proofs for Sections 2-5

Proof of Lemma 1. See text. □

Proof of Proposition 1. We solve a relaxed problem in which the seller chooses the H-types' trading probabilities (σ_A, σ_B) directly and may use buyer-specific prices $p_{1,i}$ and continuation rules ϕ_i, and then show that its solution is implementable in the commitment benchmark.

Relaxed problem. Lemma 1 does not rely on the seller's lack of commitment, so L-types never trade in the first period. The seller's revenue from buyer i is then

$$R_i = \pi_1 \Big\{ \sigma_i p_{1,i} + (1-\sigma_i)\delta \sum_z \rho_H(z)[\phi_i(z) v_H + (1-\phi_i(z)) v_L] \Big\} + (1-\pi_1)\delta v_L \sum_z \rho_L(z)[1-\phi_i(z)], \tag{A.1}$$

which is increasing in $p_{1,i}$ whenever $\sigma_i > 0$, so the seller charges each H-type her maximal acceptable price,

$$p^\star_{1,i} = v_H - \delta(v_H - v_L) \sum_z \rho_H(z)[1-\phi_i(z)]. \tag{A.2}$$

[27] Proposition B.2 (Appendix B.5) shows that the same strategic complementarities support other equilibria in which high types trade less aggressively in the first period and, as in the baseline model, the seller's revenue remains bounded away from its commitment level even as $\lambda \to 1$.

Substituting (A.2) into (A.1) and collecting terms,

$$R_i = \pi_1 v_H[\delta + (1-\delta)\sigma_i] + \delta \sum_z [1 - \phi_i(z)]\, w_i(z), \tag{A.3}$$

where $w_i(z) \equiv (1-\pi_1)v_L\rho_L(z) - \pi_1(v_H - v_L)\rho_H(z)$. Because $\lambda v_H > v_L$ from Assumption 1, we have $w_i(z_g) = \pi_1\sigma_j(v_L - \lambda v_H) \leq 0$, so $\phi_i(z_g) = 1$ is optimal, while $\phi_i(z_b) = 1$ is optimal if $w_i(z_b) < 0$ and $\phi_i(z_b) = 0$ if $w_i(z_b) > 0$. Hence

$$\max_{\phi_i(z_b)} R_i = \pi_1 v_H[\delta + (1-\delta)\sigma_i] + \delta \max\{w_i(z_b), 0\},$$

which is strictly increasing in σ_i and, as $w_i(z_b)$ increases in σ_j,[28] nondecreasing in σ_j. Thus $R_A + R_B$ is uniquely maximized at $\sigma_A = \sigma_B = 1$, where $w_i(z_g) < 0$ and

$$w_i(z_b) = (1-\pi_1)v_L - \pi_1(1-\lambda)v_H > 0 \quad \Leftrightarrow \quad \lambda > \lambda^c = 1 - \frac{(1-\pi_1)\pi^\star}{\pi_1}.$$

Note that $\lambda^c > \pi_1$ from Assumption 1 and $\lambda^c < 1$. For $\lambda \neq \lambda^c$ the relaxed problem has a unique solution: $\sigma_i = 1$, $\phi_i(z_g) = 1$, and, by (A.2), either $\phi_i(z_b) = 1$ and $p_{1,i} = v_H$ (if $\lambda < \lambda^c$) or $\phi_i(z_b) = 0$ and $p_{1,i} = v_H - \delta(1-\lambda)(v_H - v_L)$ (if $\lambda > \lambda^c$). This is the allocation in the statement.

Implementation. The relaxed solution is symmetric, so it is an admissible plan (p_1, ϕ). Under it, p_1 is the H-type's maximal acceptable price when the other H-type accepts, so immediate acceptance by both H-types is an equilibrium of the buyers' game, and it is the one the seller selects. The plan thus attains the value of the relaxed problem, which bounds the seller's commitment payoff from above, and is optimal. Since any optimal plan must attain this value, uniqueness for $\lambda \neq \lambda^c$ follows from uniqueness of the relaxed solution. □

Proof of Proposition 2. The nonmonotonicity of the inverse demand schedule is discussed in the text. By (9) and (10), $\pi_2(z_g; \sigma) = \lambda(1-\sigma)/(1-\lambda\sigma)$, so $\bar{\sigma} = (\lambda - \pi^\star)/[\lambda(1-\pi^\star)]$, which increases in λ. Also, $1/\pi_2(z_b;\sigma) - 1 = [1 - \pi_1 - \pi_1(1-\lambda)\sigma]/[\pi_1(1-\sigma)(1-\lambda\sigma)]$ increases in λ, and in σ since $\lambda \geq \pi_1$. As $\pi_2(z_b;\sigma)$ falls from $\pi_1 > \pi^\star$ at $\sigma = 0$ to zero as $\sigma \to 1$, $\underline{\sigma}$ is unique and decreases in λ. □

Proof of Proposition 3. See text. □

By Proposition 3, every equilibrium is a high-price, a low-price, or a contingent-price equilibrium, so proving Theorem 1 amounts to determining which of these candidates survive the seller's deviations at each λ. We first compute the seller's revenue in each candidate and from every off-path offer. Next, Lemmas A.1 and A.2 reduce the existence of the high-price and low-price equilibria to a ranking of these revenues. Subsequently, Lemma A.3 shows that, at any given price $p_1 \in [\underline{p}, \bar{p}]$, the seller earns weakly more when H-types trade with probability $\bar{\sigma}$ than with probability $\sigma^\star(p_1)$. Thus, the existence of a contingent-price equilibrium hinges on a single comparison, between the revenue of the best contingent-price candidate and those of the two noncontingent ones (Lemma A.4).

[28] Indeed, $w_i(z_b) = v_L - \pi_1 v_H + \pi_1\sigma_j(\lambda v_H - v_L)$.

Per-buyer revenues (Equation (17)) in the high-price and low-price equilibria are

$$R^{hp} = (1-\delta)\pi_1 \underline{\sigma} v_H + \delta \pi_1 v_H, \tag{A.4}$$

$$R^{\ell p} = (1-\delta)\pi_1 v_H + \delta v_L. \tag{A.5}$$

For $p_1 \in [\underline{p}, \bar{p}]$, contingent-price revenues with $\sigma = \bar{\sigma}$ and $\sigma = \sigma^\star(p_1)$ are

$$R^{cp}(p_1, \bar{\sigma}) = \pi_1 \bar{\sigma} p_1 + \delta(1 - \pi_1 \bar{\sigma}) v_L, \tag{A.6}$$

$$R^{cp}(p_1, \sigma^\star(p_1)) = (1-\delta)\pi_1 \sigma^\star(p_1) v_H + \delta[\pi_1 \sigma^\star(p_1) \lambda v_H + (1 - \pi_1 \sigma^\star(p_1)) v_L]. \tag{A.7}$$

In (A.6), the seller is indifferent after good news and charges v_L after bad news, so he earns v_L from every buyer who has not traded. In (A.7), he charges v_H after good news and v_L after bad news, and, by (16),

$$p_1 = v_H - \delta(v_H - v_L)\rho_H(z_b; \sigma^\star(p_1)). \tag{A.8}$$

Consider next a deviation to an off-path offer $\tilde{p}_1$. When several trading probabilities are consistent with the buyers' incentives following $\tilde{p}_1$, we select the one that yields the seller the lowest revenue. This is without loss of generality: any outcome that can be sustained by some selection of continuation play after deviations can also be sustained by the selection that is worst for the seller.

(a) If $\tilde{p}_1 \in (\bar{p}, v_H)$, H-types trade with probability $\underline{\sigma}$. The seller charges v_H after good news and is indifferent after bad news, so he earns $\pi_1 \underline{\sigma} \tilde{p}_1 + \delta \pi_1 (1 - \underline{\sigma}) v_H$, which is below R^{hp} and converges to it as $\tilde{p}_1 \uparrow v_H$. If $\tilde{p}_1 \in [\underline{p}, \bar{p}]$, trading with probability $\underline{\sigma}$ remains consistent, so the selected outcome yields at most this expression. An offer of v_H induces some $\sigma \leq \underline{\sigma}$ and yields $\pi_1 v_H[\delta + (1-\delta)\sigma] \leq R^{hp}$; higher offers are rejected and yield $\delta \pi_1 v_H < R^{hp}$.

(b) If $\tilde{p}_1 \in (v_H - \delta(v_H - v_L), \underline{p})$, H-types trade with probability $\bar{\sigma}$. The seller is indifferent after good news and charges v_L after bad news, so he earns $\pi_1 \bar{\sigma} \tilde{p}_1 + \delta(1 - \pi_1 \bar{\sigma}) v_L$, which is below $R^{cp}(\underline{p}, \bar{\sigma})$ and converges to it as $\tilde{p}_1 \uparrow \underline{p}$.

(c) If $\tilde{p}_1 \in (v_L, v_H - \delta(v_H - v_L))$, H-types accept for sure and the seller earns $\pi_1 \tilde{p}_1 + \delta(1 - \pi_1) v_L$, which is below $R^{\ell p}$ and converges to it as $\tilde{p}_1 \uparrow v_H - \delta(v_H - v_L)$. Offers of at most v_L yield at most $v_L < R^{\ell p}$. An offer of $v_H - \delta(v_H - v_L)$ induces some $\sigma \geq \bar{\sigma}$ and yields $\pi_1 \sigma \tilde{p}_1 + \delta(1 - \pi_1 \sigma) v_L \leq R^{\ell p}$.

Hence, whatever the on-path offer p_1, the supremum of the seller's revenue over offers $\tilde{p}_1 \neq p_1$ is

$$\max\{R^{hp}, R^{\ell p}, R^{cp}(\underline{p}, \bar{\sigma})\}. \tag{A.9}$$

Lemma A.1. *The high-price equilibrium exists if and only if* $R^{hp} \geq \max\{R^{\ell p}, R^{cp}(\underline{p}, \bar{\sigma})\}$.

Lemma A.2. *The low-price equilibrium exists if and only if* $R^{\ell p} \geq \max\{R^{hp}, R^{cp}(\underline{p}, \bar{\sigma})\}$.

Proof of Lemmas A.1 and A.2. The seller has no profitable deviation from the high-price (low-price) offer if and only if R^{hp} ($R^{\ell p}$) is at least (A.9). □

Lemma A.3. *For every $p_1 \in [\underline{p}, \bar{p}]$, $R^{cp}(p_1, \sigma^\star(p_1)) \leq R^{cp}(p_1, \bar{\sigma})$, with strict inequality whenever $p_1 < \bar{p}$.*

Proof. Since $\pi_2(z_g; \bar{\sigma}) = \pi^\star$, we have $\lambda v_H - v_L = \lambda\bar{\sigma}(v_H - v_L)$. Substituting (A.8) into (A.6), and this identity into (A.7), gives

$$R^{cp}(p_1, \bar{\sigma}) = (1-\delta)\pi_1\bar{\sigma}v_H + \delta[\pi_1\lambda\bar{\sigma}\sigma^\star(p_1)(v_H - v_L) + v_L],$$
$$R^{cp}(p_1, \sigma^\star(p_1)) = (1-\delta)\pi_1\sigma^\star(p_1)v_H + \delta[\pi_1\lambda\bar{\sigma}\sigma^\star(p_1)(v_H - v_L) + v_L].$$

The difference is $(1-\delta)\pi_1 v_H(\bar{\sigma} - \sigma^\star(p_1)) \geq 0$, strictly so if $p_1 < \bar{p}$, because $\sigma^\star$ is increasing with $\sigma^\star(\bar{p}) = \bar{\sigma}$. □

Lemma A.4. *A contingent-price equilibrium exists if and only if $R^{cp}(\bar{p}, \bar{\sigma}) \geq \max\{R^{hp}, R^{\ell p}\}$. Moreover, if the inequality is strict, then there exists a continuum of contingent-price equilibria.*

Proof. By Proposition 3, a contingent-price equilibrium has $p_1 \in [\underline{p}, \bar{p}]$ and $\sigma \in \{\bar{\sigma}, \sigma^\star(p_1)\}$. Since $R^{cp}(\cdot, \bar{\sigma})$ is strictly increasing in p_1, Lemma A.3 gives

$$R^{cp}(p_1, \sigma^\star(p_1)) \leq R^{cp}(p_1, \bar{\sigma}) \leq R^{cp}(\bar{p}, \bar{\sigma}) \quad \text{and} \quad R^{cp}(\underline{p}, \bar{\sigma}) \leq R^{cp}(p_1, \bar{\sigma}). \tag{A.10}$$

If $R^{cp}(\bar{p}, \bar{\sigma}) < \max\{R^{hp}, R^{\ell p}\}$, every candidate earns less than (A.9), so the seller has a profitable deviation. Conversely, by (A.9) and (A.10), $(p_1, \bar{\sigma})$ is an equilibrium if and only if $R^{cp}(p_1, \bar{\sigma}) \geq \max\{R^{hp}, R^{\ell p}\}$. This holds at $p_1 = \bar{p}$ under the stated condition and, because $R^{cp}(\cdot, \bar{\sigma})$ is continuous and strictly increasing, the set of such prices is an interval $[\hat{p}, \bar{p}]$, where $\hat{p} = \underline{p}$ if $R^{cp}(\underline{p}, \bar{\sigma}) \geq \max\{R^{hp}, R^{\ell p}\}$ and $\hat{p}$ solves $R^{cp}(\hat{p}, \bar{\sigma}) = \max\{R^{hp}, R^{\ell p}\}$ otherwise.

If $R^{cp}(\bar{p}, \bar{\sigma}) > \max\{R^{hp}, R^{\ell p}\}$, then $\hat{p} < \bar{p}$. Each $p_1 \in [\hat{p}, \bar{p}]$ is sustained with $\sigma = \bar{\sigma}$ by the mixing probability $\phi(z_g)$ that keeps H-types indifferent, that is, by the $\phi(z_g)$ that solves:

$$p_1 = v_H - \delta(v_H - v_L)[\rho_H(z_g; \bar{\sigma})(1 - \phi(z_g)) + \rho_H(z_b; \bar{\sigma})].$$

Since $\rho_H(z_g; \sigma) = \lambda\sigma = 1 - \rho_H(z_b; \sigma)$, comparing the above equation with(A.8) gives $\phi(z_g) = \sigma^\star(p_1)/\bar{\sigma}$. Hence, when $\hat{p} < \bar{p}$, there is a continuum of contingent-price equilibria, indexed by $\phi(z_g) \in [\sigma^\star(\hat{p})/\bar{\sigma}, 1]$, which equals $[\underline{\sigma}/\bar{\sigma}, 1]$ if $\hat{p} = \underline{p}$. □

Having established Lemmas A.1 to A.4, we proceed to proving Theorem 1.

Proof of Theorem 1. Let $\Delta^{hp} \equiv R^{cp}(\bar{p}, \bar{\sigma}) - R^{hp}$, $\Delta^{\ell p} \equiv R^{cp}(\bar{p}, \bar{\sigma}) - R^{\ell p}$ and $\Delta \equiv \min\{\Delta^{hp}, \Delta^{\ell p}\}$. We first claim there is a unique $\lambda^{nc} \in (\pi_1, 1)$, depending on δ, with $\Delta < 0$ for $\lambda < \lambda^{nc}$ and $\Delta > 0$ for $\lambda > \lambda^{nc}$. Indeed, since $\bar{p} = v_H - \delta(1 - \lambda\bar{\sigma})(v_H - v_L)$,

$$\Delta^{\ell p}(\lambda) = \pi_1 v_H \left[-(1-\delta)(1-\bar{\sigma}) + \delta\lambda\bar{\sigma}^2(1-\pi^\star)\right], \tag{A.11}$$
$$\Delta^{hp}(\lambda) = \pi_1 v_H(1-\delta)(\bar{\sigma} - \underline{\sigma}) + \delta v_H \left[\pi_1\lambda\bar{\sigma}^2(1-\pi^\star) - (\pi_1 - \pi^\star)\right]. \tag{A.12}$$

Both $\bar{\sigma}$ and $\underline{\sigma}$ are continuous in λ and, by Proposition 2, $\bar{\sigma}$ is strictly increasing and $\underline{\sigma}$ strictly decreasing. Hence $\Delta^{\ell p}$, Δ^{hp}, and Δ are continuous and strictly increasing in λ. At $\lambda = \pi_1$, $\underline{\sigma} = \bar{\sigma} = \sigma^0 \equiv (\pi_1 - \pi^\star)/[\pi_1(1-\pi^\star)]$, so

$$\Delta(\pi_1) \leq \Delta^{hp}(\pi_1) = -\delta v_H \frac{(\pi_1 - \pi^\star)(1-\pi_1)}{1-\pi^\star} < 0.$$

As $\lambda \to 1$, $\bar{\sigma} \to 1$, so $\Delta^{\ell p} \to \delta\pi_1(v_H - v_L) > 0$; and since $\bar{\sigma} \geq \underline{\sigma}$, Δ^{hp} is bounded below by the second term of (A.12), which converges to $\delta(1-\pi_1)v_L > 0$. Hence there is a unique $\lambda^{nc} \in (\pi_1, 1)$, depending on δ, with $\Delta < 0$ for $\lambda < \lambda^{nc}$ and $\Delta > 0$ for $\lambda > \lambda^{nc}$. We write $\lambda^{nc}(\delta)$ where the dependence matters.

Next, we prove part (i) of the theorem. Let $\lambda < \lambda^{nc}$, so $R^{cp}(\underline{p}, \bar{\sigma}) < R^{cp}(\bar{p}, \bar{\sigma}) < \max\{R^{hp}, R^{\ell p}\}$. By Lemma A.4, no contingent-price equilibrium exists. By Lemmas A.1 and A.2, the high-price equilibrium exists if and only if $R^{hp} \geq R^{\ell p}$, and the low-price equilibrium if and only if $R^{\ell p} \geq R^{hp}$. The equilibrium is therefore unique unless $R^{hp} = R^{\ell p}$.

Now, we prove part (ii) of the theorem. Let $\lambda > \lambda^{nc}$. Then $R^{cp}(\bar{p}, \bar{\sigma}) > \max\{R^{hp}, R^{\ell p}\}$ and Lemma A.4 yields a continuum of contingent-price equilibria. To see that either noncontingent equilibrium may coexist with them, note that

$$R^{hp} - R^{\ell p} = v_H \left\{\delta(\pi_1 - \pi^\star) - (1-\delta)\pi_1(1-\underline{\sigma})\right\},$$

where $1 - \underline{\sigma} \in [1 - \sigma^0, 1]$ for every λ, because $\underline{\sigma}$ is decreasing in λ. Hence $R^{\ell p} > R^{hp}$ for all λ if δ is sufficiently low, and $R^{hp} > R^{\ell p}$ for all λ if δ is sufficiently high. In the first case, $\Delta = 0$ at $\lambda = \lambda^{nc}$ gives $R^{\ell p} = R^{cp}(\bar{p}, \bar{\sigma}) > R^{cp}(\underline{p}, \bar{\sigma})$. By continuity in λ, $R^{\ell p} > \max\{R^{hp}, R^{cp}(\underline{p}, \bar{\sigma})\}$ on $(\lambda^{nc}, \lambda^{nc} + \varepsilon)$ for some $\varepsilon > 0$, where by Lemma A.2 the low-price equilibrium coexists with the continuum. In the second case, the same argument with hp and ℓp interchanged gives coexistence with the high-price equilibrium. The contingent-price equilibria coexist with either the low-price or the high-price equilibrium, because $R^{hp} \neq R^{\ell p}$ except in knife-edge cases.

Finally, we establish the existence of equilibria for all λ. Parts (i) and (ii) above cover $\lambda \neq \lambda^{nc}$. At $\lambda = \lambda^{nc}$, $(\bar{p}, \bar{\sigma})$ is an equilibrium by Lemma A.4. □

Proof of Proposition 4. In every equilibrium, L-types earn zero and H-types, who are willing to accept p_1, earn $v_H - p_1$. Hence one equilibrium Pareto dominates another if and only if it has a weakly lower first-period offer and a weakly higher revenue, with at least one strict inequality.

Since, generically, multiplicity only happens generically when contingent-price equilibria exist, let $\lambda > \lambda^{nc}$, so that $R^{cp}(\bar{p}, \bar{\sigma}) > \max\{R^{hp}, R^{\ell p}\}$. The high-price equilibrium, whenever it exists, is dominated by $(\bar{p}, \bar{\sigma})$, which has a lower offer and a higher revenue. A contingent-price equilibrium $(p_1, \sigma^\star(p_1))$ with $p_1 < \bar{p}$, whenever it exists, is dominated by $(p_1, \bar{\sigma})$. Indeed, since $(p_1, \sigma^\star(p_1))$ is an equilibrium, its revenue is at least (A.9). By Lemma A.3, $(p_1, \bar{\sigma})$ yields strictly more, so the seller has no profitable deviation from it either. Hence $(p_1, \bar{\sigma})$ is an equilibrium with the same offer and a strictly higher revenue.

Turning to undominated equilibria, it remains to compare the low-price equilibrium and the contingent-price equilibria with $\sigma = \bar{\sigma}$. Among the latter, a higher offer yields a higher revenue, as $R^{cp}(\cdot, \bar{\sigma})$ is strictly increasing, so none dominates another. The low-price equilibrium has the lowest offer, so it is undominated.

It dominates $(p_1, \bar{\sigma})$ if and only if $R^{\ell p} \geq R^{cp}(p_1, \bar{\sigma})$; since every equilibrium revenue is at least (A.9), this requires $R^{cp}(p_1, \bar{\sigma}) = R^{\ell p}$, which holds only at the lowest offer among these equilibria. □

Proof of Proposition 5. Fix $\delta \in (0, 1)$ and make the dependence of revenues on λ explicit. We first characterize the lowest and highest attainable revenues:

$$\underline{R}(\lambda) = \max\left\{R^{hp}(\lambda), R^{\ell p}, R^{cp}(\underline{p}, \bar{\sigma}; \lambda)\right\}, \tag{A.13}$$

$$\overline{R}(\lambda) = \max\left\{R^{hp}(\lambda), R^{\ell p}, R^{cp}(\bar{p}, \bar{\sigma}; \lambda)\right\}. \tag{A.14}$$

By (A.9), every equilibrium revenue is at least the right-hand side of (A.13). This bound is attained: by the high-price or the low-price equilibrium if the maximum is $R^{hp}(\lambda)$ or $R^{\ell p}$ (Lemmas A.1 and A.2), and otherwise by $(\underline{p}, \bar{\sigma})$, which is an equilibrium when $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) > \max\{R^{hp}(\lambda), R^{\ell p}\}$ (proof of Lemma A.4).

As for (A.14), by Proposition 3 and (A.10), no equilibrium yields more than its right-hand side. This bound is attained by $(\bar{p}, \bar{\sigma})$ if $R^{cp}(\bar{p}, \bar{\sigma}; \lambda) \geq \max\{R^{hp}(\lambda), R^{\ell p}\}$ (Lemma A.4) and otherwise, that is, for $\lambda < \lambda^{nc}$, by the high-price or the low-price equilibrium (Theorem 1).

We next show that $\mathcal{R}(\lambda)$ is a closed interval, and that its endpoints are continuous, which gives hemi-continuity. By definition, $\mathcal{R}(\lambda) \subseteq [\underline{R}(\lambda), \overline{R}(\lambda)]$. If $\lambda < \lambda^{nc}$, then $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) < R^{cp}(\bar{p}, \bar{\sigma}; \lambda) < \max\{R^{hp}(\lambda), R^{\ell p}\}$, so $\underline{R}(\lambda) = \overline{R}(\lambda)$ and $\mathcal{R}(\lambda)$ is a singleton. If $\lambda \geq \lambda^{nc}$, recall that $(p_1, \bar{\sigma})$ is an equilibrium if and only if $R^{cp}(p_1, \bar{\sigma}; \lambda) \geq \max\{R^{hp}(\lambda), R^{\ell p}\}$. Since $R^{cp}(\cdot, \bar{\sigma}; \lambda)$ is continuous and strictly increasing on $[\underline{p}, \bar{p}]$, these equilibria attain every revenue between $\underline{R}(\lambda)$ and $\overline{R}(\lambda)$. Hence, for every λ,

$$\mathcal{R}(\lambda) = [\underline{R}(\lambda), \overline{R}(\lambda)]. \tag{A.15}$$

Moreover, $\underline{\sigma}$ and $\bar{\sigma}$ are continuous in λ, and hence so are $\underline{p}$, $\bar{p}$ and the revenues in (A.13) and (A.14). Thus $\underline{R}$ and $\overline{R}$ are continuous, and a compact-interval-valued correspondence with continuous endpoints is both upper and lower hemi-continuous.

We now turn to the comparative statics of the two endpoints. For part (1), let $\lambda < \lambda^{nc}$, so that $\underline{R}(\lambda) = \overline{R}(\lambda) = \max\{R^{hp}(\lambda), R^{\ell p}\}$. This is weakly decreasing, because $R^{hp}(\lambda)$ decreases with λ, as $\underline{\sigma}$ does, and $R^{\ell p}$ does not depend on λ.

For part (2.i), we first show that $\underline{R}$ can be nonmonotone. By (A.13), $\underline{R}$ is the upper envelope of $R^{hp}(\lambda)$, which is strictly decreasing, of the constant $R^{\ell p}$, and of $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) = \delta v_L + \pi_1 \bar{\sigma}(\underline{p} - \delta v_L)$, which may increase with λ because $\bar{\sigma}$ does. It is therefore nonmonotone whenever R^{hp} determines it for low λ and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$ determines it for high λ. For instance, if $\pi_1 = 0.7$, $v_L = 1$, $v_H = 2$, and $\delta = 0.645$, then $\lambda^{nc} \approx 0.763$ and $\underline{R}$ equals R^{hp} on $[\lambda^{nc}, 0.782]$, $R^{\ell p}$ on $[0.782, 0.798]$, and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$, which is strictly increasing, on $[0.798, 1)$. Figure 4 illustrates this parameterization.

We next show that $\underline{R}$ remains bounded away from the seller's revenue under commitment, $R^c(\lambda)$. By the proof of Proposition 1, any outcome in which H-types trade in the first period with probability below one yields strictly less than $R^c(\lambda)$. Since $\underline{\sigma}, \bar{\sigma} < 1$, it follows that $R^{hp}(\lambda) < R^c(\lambda)$ and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) < R^c(\lambda)$. Moreover, $R^{\ell p} < \pi_1 v_H \leq R^c(\lambda)$. Hence $\underline{R}(\lambda) < R^c(\lambda)$ for every $\lambda < 1$ and, by continuity, it suffices to

show that the gap does not vanish as $\lambda \to 1$. Let $\underline{\sigma}(1) \equiv \lim_{\lambda\to 1} \underline{\sigma}(\lambda)$, which is below one because $\underline{\sigma}$ is decreasing. Since $\bar{\sigma} \to 1$ and $\underline{p} \to v_H - \delta(1 - \underline{\sigma}(1))(v_H - v_L)$,

$$\lim_{\lambda\to 1} R^{hp}(\lambda) = \pi_1 v_H \left[\delta + (1-\delta)\underline{\sigma}(1)\right] < \pi_1 v_H,$$
$$\lim_{\lambda\to 1} R^{cp}(\underline{p}, \bar{\sigma}; \lambda) = R^c(1) - \pi_1 \delta(1 - \underline{\sigma}(1))(v_H - v_L),$$

where $R^c(1) = \pi_1 v_H + (1-\pi_1)\delta v_L > \pi_1 v_H > R^{\ell p}$. All three limits are strictly below $R^c(1)$.

Finally, for part (2.ii), let $\lambda \geq \lambda^{nc}$. Then $R^{cp}(\bar{p}, \bar{\sigma}; \lambda) \geq \max\{R^{hp}(\lambda), R^{\ell p}\}$, so (A.14) gives $\overline{R}(\lambda) = R^{cp}(\bar{p}, \bar{\sigma}; \lambda) = \delta v_L + \pi_1 \bar{\sigma}(\bar{p} - \delta v_L)$. Both $\bar{\sigma}$ and $\bar{p}$ are strictly increasing in λ, and $\bar{p} > \delta v_L$, so $\overline{R}$ is strictly increasing. As $\lambda \to 1$, $\bar{\sigma} \to 1$ and $\bar{p} \to v_H$, so $\overline{R}(\lambda) \to \pi_1 v_H + (1-\pi_1)\delta v_L = R^c(1)$. □

Proof of Proposition 6. Fix $\delta \in (0,1)$. We first show that $\mathcal{W}$ is not lower hemi-continuous at $\lambda = \lambda^{nc}$, where $R^{cp}(\bar{p}, \bar{\sigma}) = \max\{R^{hp}, R^{\ell p}\}$. By (A.10), every other contingent-price candidate earns strictly less, so $(\bar{p}, \bar{\sigma})$ is the only contingent-price equilibrium. By Lemmas A.1 and A.2, the equilibrium $j \in \{hp, \ell p\}$ exists if and only if $R^j = R^{cp}(\bar{p}, \bar{\sigma})$, which holds for at least one j. Such an equilibrium has the same revenue as $(\bar{p}, \bar{\sigma})$ but a different first-period offer p_1^j, so (18) gives

$$W^{cp}(\bar{p}, \bar{\sigma}) - W^j = \pi_1(p_1^j - \bar{p}) \neq 0.$$

Lower hemi-continuity of $\mathcal{W}$ fails at λ^{nc} because $W^{cp}(\bar{p}, \bar{\sigma})$ cannot be approached by equilibrium surpluses from below as $\lambda \uparrow \lambda^{nc}$. Indeed, for $\lambda < \lambda^{nc}$ every equilibrium is high-price or low-price, and the equilibrium j exists only if $R^j = \max\{R^{hp}, R^{\ell p}\}$. By continuity, as $\lambda \uparrow \lambda^{nc}$ the surplus of such equilibria can only converge to W^j for some j with $R^j = R^{cp}(\bar{p}, \bar{\sigma})$ at λ^{nc}, which differs from $W^{cp}(\bar{p}, \bar{\sigma})$ by the above equation.

To prove parts (i)–(iii) of the proposition, it is useful to rewrite the social surplus in terms of the allocation. In any equilibrium, H-types are indifferent between accepting p_1 and waiting, that is, $v_H - p_1 = \delta(v_H - v_L)\sum_z \rho_H(z;\sigma)(1 - \phi(z))$. Substituting this condition and (17) into (18) gives

$$W(\sigma, \phi; \lambda) = \pi_1[\sigma + (1-\sigma)\delta]v_H + (1-\pi_1)\delta \sum_{z\in\{z_b, z_g\}} \rho_L(z;\sigma)(1-\phi(z))v_L. \tag{A.16}$$

Hence social surplus is at most $\pi_1 v_H + (1-\pi_1)\delta v_L$ in every equilibrium, with equality in the low-price one. This bound is also the limit of the social surplus under commitment, $W^c(\lambda)$. Indeed, by Proposition 1, for $\lambda > \lambda^c$ the commitment outcome has $\sigma = 1$, $\phi(z_g) = 1$ and $\phi(z_b) = 0$, so (A.16) gives $W^c(\lambda) = \pi_1 v_H + \delta v_L[1 - \pi_1 - \pi_1(1-\lambda)]$, which converges to $W^c(1) = \pi_1 v_H + (1-\pi_1)\delta v_L$ as $\lambda \to 1$.

To prove part (i), we exhibit an equilibrium whose surplus stays bounded away from $W^c(1)$. Let $\lambda > \lambda^{nc}$ and consider the contingent-price candidates $(p_1, \sigma^\star(p_1))$. By (A.9) and (A.13), such a candidate is an equilibrium if and only if $R^{cp}(p_1, \sigma^\star(p_1); \lambda) \geq \underline{R}(\lambda)$. By (A.7), this revenue is continuous and strictly increasing in p_1, as $\sigma^\star$ is. At $p_1 = \underline{p}$, where $\sigma^\star(\underline{p}) = \underline{\sigma}$, it equals $R^{hp}(\lambda) - \pi_1\underline{\sigma}(v_H - \underline{p}) < \underline{R}(\lambda)$, and at $p_1 = \bar{p}$ it equals $\overline{R}(\lambda) > \underline{R}(\lambda)$. Hence there is a unique $\tilde{p}_1$ such that $R^{cp}(\tilde{p}_1, \sigma^\star(\tilde{p}_1); \lambda) = \underline{R}(\lambda)$, and

$(\tilde{p}_1, \sigma^\star(\tilde{p}_1))$ is an equilibrium. In it, $\phi(z_g) = 1$ and $\phi(z_b) = 0$, so (A.7) and (A.16) give

$$R^c(1) - \underline{R}(\lambda) = \pi_1[1 - \sigma^\star(\tilde{p}_1)](v_H - \delta v_L) + \pi_1\delta\sigma^\star(\tilde{p}_1)(1-\lambda)v_H,$$
$$W^c(1) - W^{cp}(\tilde{p}_1, \sigma^\star(\tilde{p}_1); \lambda) = \pi_1[1 - \sigma^\star(\tilde{p}_1)](1-\delta)v_H + \pi_1\delta\sigma^\star(\tilde{p}_1)(1-\lambda)v_L.$$

By Proposition 5, the left-hand side of the first line remains bounded away from zero as $\lambda \to 1$. Since the last term of that line vanishes, $1 - \sigma^\star(\tilde{p}_1)$ remains bounded away from zero, and therefore so does the second line. Part (i) follows because $\underline{W}(\lambda) \leq W^{cp}(\tilde{p}_1, \sigma^\star(\tilde{p}_1); \lambda)$.

Finally, we prove parts (ii) and (iii). By Proposition 4, a Pareto-efficient equilibrium is either the low-price equilibrium, whose surplus is $W^c(1)$, or a contingent-price equilibrium with $\sigma = \bar{\sigma}$. In the latter, L-types trade at least after bad news, so by (A.16) its surplus is at least

$$\pi_1[\bar{\sigma} + (1-\bar{\sigma})\delta]v_H + \delta v_L[1 - \pi_1 - \pi_1(1-\lambda)\bar{\sigma}],$$

which converges to $W^c(1)$ as $\lambda \to 1$, because $\bar{\sigma} \to 1$. Since no equilibrium surplus exceeds $W^c(1)$, this proves part (iii). Part (ii) follows, because $\overline{W}(\lambda)$ is at least the surplus of a Pareto-efficient equilibrium and at most $W^c(1)$. □

Proof of Proposition 7. See text. □

Proof of Proposition 8. See text. □

Proof of Corollary 1. Without investment, $\lambda = \pi_1 < \lambda^{nc}$, and the unique equilibrium yields revenue $\overline{R}(\pi_1)$. Let R denote the revenue in the continuation equilibrium played after investment, so that investment is profitable if $2R - c > 2\overline{R}(\pi_1)$. Since $R \leq \overline{R}(\lambda^D)$, profitability requires $2\overline{R}(\lambda^D) - c > 2\overline{R}(\pi_1)$. For $\lambda \leq \lambda^{nc}$, $\overline{R}(\lambda) \leq \overline{R}(\pi_1)$; for $\lambda > \lambda^{nc}$, $\overline{R}(\lambda)$ is strictly increasing in λ, by Proposition 5. Hence the seller finds investment profitable only if $\lambda^D > \bar{\lambda}^D$, where $\bar{\lambda}^D > \lambda^{nc}$ solves $2\overline{R}(\bar{\lambda}^D) - c = 2\overline{R}(\pi_1)$, with $\bar{\lambda}^D = 1$ if no solution exists.

If the equilibrium without investment is the low-price one, it attains the maximal social surplus $\pi_1 v_H + (1-\pi_1)\delta v_L$ per buyer; profitable investment thus lowers social surplus and, since revenue rises, buyer surplus. If it is the high-price one, buyer surplus is zero and social surplus equals $2\overline{R}(\pi_1)$, so profitable investment weakly raises buyer surplus, strictly in any contingent-price equilibrium, and raises social surplus net of c, which is at least $2R - c > 2\overline{R}(\pi_1)$. In either case, profitable investment exists for λ^D close enough to one whenever $c < 2\left(W^c(1) - \overline{R}(\pi_1)\right)$, since $\overline{R}(\lambda) \to W^c(1)$ as $\lambda \to 1$. □

Proof of Corollary 2. Under the experimentation policy, buyer A is offered $p_1^A = v_H - \delta(v_H - v_L)$ and therefore reveals her type through her first-period trading decision. Buyer B does not trade in the first period. Following trade by A, the seller's posterior that B is an H-type is $\lambda > \pi^\star$. Following no trade, it is $\pi_1(1-\lambda)/(1-\pi_1)$. Hence, recalling the proof of Proposition 1, buyer B's second-period price is strictly contingent on buyer A's trading decision if and only if

$$\lambda > \lambda^c = 1 - \frac{(1-\pi_1)\pi^\star}{\pi_1}.$$

For $\lambda < \lambda^c$, the seller therefore charges v_H to buyer B following either realization of buyer A's trading decision. Experimentation then yields total revenue:

$$R^{A/B} = R^{\ell p} + \delta \pi_1 v_H.$$

Since $R^{hp}(\lambda) = \delta \pi_1 v_H + (1-\delta)\pi_1 \underline{\sigma}(\lambda) v_H \geq \delta \pi_1 v_H$, we have:

$$\begin{aligned} R^{A/B} &\leq R^{\ell p} + R^{hp}(\lambda), \\ &\leq 2 \max \left\{ R^{\ell p}, R^{hp}(\lambda) \right\}, \\ &\leq 2 \max \left\{ R^{\ell p}, R^{hp}(\lambda), R^{cp}(\underline{p}, \bar{\sigma}; \lambda) \right\}, \\ &= 2\underline{R}(\lambda). \end{aligned}$$

Thus, experimentation cannot be strictly profitable unless $\lambda > \lambda^c$. For $\lambda > \lambda^c$, the seller charges v_H to buyer B following trade by buyer A and v_L following no trade. Experimentation then yields

$$R^{A/B}(\lambda) = R^{\ell p} + \delta \left[\pi_1 \lambda v_H + (1 - \pi_1) v_L \right].$$

Whether this revenue exceeds $2\underline{R}(\lambda)$ depends on the remaining parameters. Define:

$$\lambda^{A/B} \equiv \inf \left\{ \lambda \in (\pi_1, 1) : R^{A/B}(\lambda) > 2\underline{R}(\lambda) \right\},$$

with $\lambda^{A/B} = 1$ if this set is empty. Since both revenue functions are continuous, the set in the definition of $\lambda^{A/B}$ is open. The argument above implies $\lambda^{A/B} \geq \lambda^c$. At $\lambda = \lambda^c$ the two expressions for $R^{A/B}$ coincide, and the inequality $R^{A/B}(\lambda^c) < R^{\ell p} + R^{hp}(\lambda^c) \leq 2\underline{R}(\lambda^c)$ is strict because $\underline{\sigma} > 0$. By continuity, experimentation is unprofitable on a neighborhood of λ^c, so $\lambda^{A/B} > \lambda^c$. Hence, whenever experimentation is profitable, necessarily $\lambda > \lambda^{A/B}$.

Finally, buyer surplus and social surplus may either increase or decrease. Under the experimentation policy with $\lambda > \lambda^c$, buyer A earns $\delta(v_H - v_L)$ if she is an H-type, and B earns $\delta(v_H - v_L)$ only if she is an H-type and A is not, so buyer surplus equals $\pi_1 \delta (v_H - v_L)(2 - \lambda)$. If the benchmark is the high-price equilibrium, as in Panel A of Figure 6, buyers earn zero without experimentation, so whenever experimentation is profitable both buyer surplus and social surplus rise. If the low-price equilibrium is played absent experimentation, as can happen in Panel B of Figure 6 for λ slightly above $\lambda^{A/B}$ (e.g., $\lambda = 0.84$), buyer surplus falls from $2\pi_1 \delta (v_H - v_L)$ to $\pi_1 \delta (v_H - v_L)(2 - \lambda)$, and social surplus falls below $2[\pi_1 v_H + (1 - \pi_1)\delta v_L]$, the maximum attained by the low-price equilibrium, because buyer B trades only in the second period. □

Proof of Corollary 3. Suppose first that $\lambda < \lambda^{nc}$ and that δ is sufficiently high. The unique continuation equilibrium is then the high-price equilibrium,[29] in which suppliers obtain zero surplus. An investing supplier obtains payoff $-F$, whereas not investing guarantees at least zero. Hence, a pure both-invest equilibrium with this continuation is impossible.

[29] Uniqueness requires $R^{hp} > R^{\ell p}$, i.e., $\delta > \frac{\pi_1(1-\underline{\sigma})}{\pi_1(1-\underline{\sigma}) + \pi_1 - \pi^\star}$, a threshold increasing in λ with supremum below one.

Now suppose $\lambda > \lambda^{nc}$. Consider a contingent-price equilibrium with $\sigma = \bar{\sigma}$ in which the downstream buyer charges αv_H following good news with probability $\phi \in [\underline{\sigma}/\bar{\sigma}, 1]$, and the first-period price is $\alpha p_1(\phi)$ with $p_1(\phi) = v_H - \delta(1 - \lambda\bar{\sigma}\phi)(v_H - v_L)$. Let $\underline{\phi}$ denote the smallest ϕ sustaining the equilibrium, i.e., satisfying $R^{cp}(p_1(\phi), \bar{\sigma}) \geq \max\left\{R^{hp}, R^{\ell p}\right\}$. Since the downstream buyer is indifferent following good news, $R^{cp}(p_1(\phi), \bar{\sigma})$ is strictly increasing in ϕ, so $\underline{\phi}$ is either $\underline{\sigma}/\bar{\sigma}$ or the unique solution of $R^{cp}(p_1(\phi), \bar{\sigma}) = \max\{R^{hp}, R^{\ell p}\}$. Since investment scales both valuations proportionally, it leaves $\pi^\star$, $\underline{\sigma}$, $\bar{\sigma}$, and $\underline{\phi}$ unchanged and scales all revenues by α, so this equilibrium exists when both suppliers invest.

If both suppliers invest and coordinate on the contingent-price equilibrium with $\sigma = \bar{\sigma}$ and $\phi(z_g) = \underline{\phi}$, each supplier obtains ex ante $-F + \alpha\delta\pi_1(v_H - v_L)\left(1 - \lambda\bar{\sigma}\underline{\phi}\right)$. Let

$$\underline{\alpha} \equiv \max\left\{\frac{2v_H}{v_L},\ \frac{F}{\delta\pi_1(v_H - v_L)\left(1 - \lambda\bar{\sigma}\underline{\phi}\right)}\right\}.$$

Note that $\underline{\alpha}$ exceeds one because of the first term and finite because of the second, since $\lambda\bar{\sigma}\underline{\phi} \leq \lambda < 1$. It moreover remains bounded as $\lambda \to 1$, since $\underline{\phi}$ remains bounded away from one. Suppose $\alpha \geq \underline{\alpha}$. By the second term, the surplus above is nonnegative. Since an equilibrium specifies continuation play following every investment profile, it therefore suffices to construct, following a unilateral deviation to no investment, a continuation equilibrium in which the deviating supplier obtains zero surplus.[30]

Call the deviating supplier N and the investing supplier I. In the continuation equilibrium we construct, N never trades in the first period, so the downstream buyer learns nothing about I from N's behavior, and I faces a single-buyer problem without information spillovers. Contingent pricing toward I is then impossible, and the relevant first-period offers are those of the single-buyer problem.

A supplier accepts a second-period offer if and only if her valuation is at least the offer, so second-period offers lie in $\{v_L, v_H, \alpha v_L, \alpha v_H\}$. If both suppliers delayed, an offer of v_L or v_H is also accepted by I and yields at most $2v_H \leq \alpha v_L$, which αv_L guarantees; hence the downstream buyer offers αv_H or αv_L, according to whether his posterior about I is above or below $\pi^\star$ (mixing if equal), and N does not trade. He follows the same rule if only I delayed, and, if only N delayed, applies the same rule to v_H and v_L according to his posterior about N. Thus I's value from delaying is zero if she is an L-type and at most $\delta\alpha(v_H - v_L)$ if she is an H-type, and N trades in the second period only if I traded in the first.

Let $\sigma^0 \in (0, 1)$ solve $\pi_1(1 - \sigma^0)/[\pi_1(1 - \sigma^0) + 1 - \pi_1] = \pi^\star$. Following a first-period offer p_1, play is as follows.

(a) $p_1 \leq v_H$: both types of I strictly accept, since $\alpha v_L \geq 2v_H$ implies $\alpha v_\theta - p_1 > \delta\alpha(v_\theta - v_L)$. I's trade is then uninformative, and N plays the standard single-buyer continuation. The downstream buyer earns at most $2v_H$.

(b) $v_H < p_1 \leq \alpha v_L$: both types of I accept and N rejects; the posterior about N is $\pi_1 > \pi^\star$, so she is offered v_H.

[30] The result does not rely on the choice of continuation: since v_L is accepted by every supplier type in the second period, no price is below v_L, so a noninvestor earns at most $\pi_1(v_H - v_L)$, which an investor's payoff exceeds once $\alpha \geq [F + \pi_1(v_H - v_L)]/[\delta\pi_1(v_H - v_L)(1 - \lambda\bar{\sigma}\underline{\phi})]$.

(c) $\alpha v_L < p_1 \leq \alpha[v_H - \delta(v_H - v_L)]$: only I's H-type accepts. Following her trade, the posterior about N is $\lambda > \pi_1 > \pi^\star$, so N is offered v_H; following no trade, I is offered αv_L.

(d) $\alpha[v_H - \delta(v_H - v_L)] < p_1 \leq \alpha v_H$: I's H-type accepts with probability σ^0. Following no trade, the posterior about I is $\pi^\star$, and the downstream buyer offers αv_H with the probability that makes her indifferent; following her trade, N is offered v_H.

(e) $p_1 > \alpha v_H$: no supplier accepts, which yields less than (d) at $p_1 = \alpha v_H$.

Within each of (b)–(d), trading probabilities and the downstream buyer's second-period revenue do not depend on p_1, so his payoff is increasing in p_1; moreover, the payoff in (a) is below that in (b) at $p_1 = \alpha v_L$. Hence he optimally offers one of αv_L, $\alpha[v_H - \delta(v_H - v_L)]$, and αv_H, and these strategies form an equilibrium of the continuation game. In it, N never trades in the first period and pays v_H whenever she trades in the second, so her surplus is zero. We conclude that, for every $\alpha \geq \underline{\alpha}$, there exists a symmetric equilibrium in which both suppliers invest.[31] □

References

Adams, Jonathan J, Min Fang, Zheng Liu, and Yajie Wang. 2025. "The rise of AI pricing: Trends, driving forces, and implications for firm performance." *Journal of Monetary Economics*, 103875.

Arieli, Itai, Moran Koren, and Rann Smorodinsky. 2022. "The implications of pricing on social learning." *Theoretical Economics*, 17(4): 1761–1802.

Asriyan, Vladimir, William Fuchs, and Brett Green. 2017. "Information spillovers in asset markets with correlated values." *American Economic Review*, 107(7): 2007–2040.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The limits of price discrimination." *American Economic Review*, 105(3): 921–957.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2025. "On the alignment of consumer surplus and total surplus under competitive price discrimination." *American Economic Journal: Microeconomics*, 17(4): 234–259.

Board, Simon, and Andrzej Skrzypacz. 2016. "Revenue management with forward-looking buyers." *Journal of Political Economy*, 124(4): 1046–1087.

Bonatti, Alessandro, and Gonzalo Cisternas. 2020. "Consumer scores and price discrimination." *Review of Economic Studies*, 87(2): 750–791.

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. 2006. "Dynamic monopoly pricing and herding." *The RAND Journal of Economics*, 37(4): 910–928.

[31] Following the profile in which neither supplier invests, which cannot be reached by a unilateral deviation, continuation play can be any equilibrium of the baseline game, which exists by Theorem 1.

Brynjolfsson, Erik, and Kristina McElheran. 2016. "The rapid adoption of data-driven decision-making." *American Economic Review*, 106(5): 133–139.

Bulow, Jeremy I. 1982. "Durable-goods monopolists." *Journal of Political Economy*, 90(2): 314–332.

Calzolari, Giacomo, Leonardo Felli, Johannes Koenen, Giancarlo Spagnolo, and Konrad O Stahl. 2021. "Relational contracts and trust in a high-tech industry." *CEPR Discussion Paper No. DP16653*.

Chamley, Christophe, and Douglas Gale. 1994. "Information revelation and strategic delay in a model of investment." *Econometrica*, 62(5): 1065–1085.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender systems as mechanisms for social learning." *The Quarterly Journal of Economics*, 133(2): 871–925.

Coase, Ronald H. 1972. "Durability and monopoly." *The Journal of Law and Economics*, 15(1): 143–149.

Crémer, Jacques, and Richard P McLean. 1988. "Full extraction of the surplus in Bayesian and dominant strategy auctions." *Econometrica*, 56(6): 1247–1257.

Drugov, Mikhail. 2021. "Bargaining with informational and payoff externalities." *Journal of Economics & Management Strategy*, 30(2): 398–419.

Dyer, Jeffrey H, and Kentaro Nobeoka. 2000. "Creating and managing a high-performance knowledge-sharing network: the Toyota case." *Strategic Management Journal*, 21(3): 345–367.

Elliott, Matthew, Andrea Galeotti, Andrew Koh, and Wenhao Li. 2026. "Market segmentation through information." *Review of Economic Studies*, rdag067.

Fuchs, William, and Andrzej Skrzypacz. 2013. "Bargaining with deadlines and private information." *American Economic Journal: Microeconomics*, 5(4): 219–243.

Fuchs, William, Aniko Öry, and Andrzej Skrzypacz. 2016. "Transparency and distressed sales under asymmetric information." *Theoretical Economics*, 11(3): 1103–1144.

Fudenberg, Drew, and Jean Tirole. 1983. "Sequential bargaining with incomplete information." *Review of Economic Studies*, 50(2): 221–247.

Fudenberg, Drew, and Jean Tirole. 1991. "Sequential bargaining under incomplete information." In *Game Theory*. Chapter 10, 397–434. MIT Press.

Fudenberg, Drew, and J Miguel Villas-Boas. 2006. "Behavior-based price discrimination and customer recognition." In *Handbook on economics and information systems*. , ed. Terrence J. Hendershott, Chapter 8, 377–436. Elsevier.

Fudenberg, Drew, David Levine, and Jean Tirole. 1985. "Infinite-horizon models of bargaining with one-sided incomplete information." In *Game-theoretic models of bargaining*. , ed. Alvin E. Roth, Chapter 5, 73–98. Cambridge University Press.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson. 1986. "Foundations of dynamic monopoly and the Coase conjecture." *Journal of Economic Theory*, 39(1): 155–190.

Hannak, Aniko, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson. 2014. "Measuring price discrimination and steering on e-commerce web sites." *Proceedings of the 2014 Conference on Internet Measurement Conference*, 305–318.

Joskow, Paul L. 1987. "Contract duration and relationship-specific investments: Empirical evidence from coal markets." *American Economic Review*, 77(1): 168–185.

Kartik, Navin, and Weijie Zhong. 2026. "Lemonade from lemons: Information design and adverse selection." *Theoretical Economics*, 21(1): 281–324.

Kohavi, Ron, and Stefan Thomke. 2017. "The surprising power of online experiments: Getting the most out of A/B and other controlled tests." *Harvard Business Review*, 95(5): 74–82.

Kohavi, Ron, Diane Tang, and Ya Xu. 2020. *Trustworthy online controlled experiments: A practical guide to A/B testing*. Cambridge University Press.

Koning, Rembrand, Sharique Hasan, and Aaron Chatterji. 2022. "Experimentation and start-up performance: Evidence from A/B testing." *Management Science*, 68(9): 6434–6453.

Kremer, Ilan, Yishay Mansour, and Motty Perry. 2014. "Implementing the 'wisdom of the crowd'." *Journal of Political Economy*, 122(5): 988–1012.

Laiho, Tuomas, and Julia Salmi. 2026. "Coasian dynamics and endogenous learning." *Journal of Economic Theory*, 106136.

Levin, Jonathan. 2003. "Relational incentive contracts." *American Economic Review*, 93(3): 835–857.

Libgober, Jonathan, and Xiaosheng Mu. 2021. "Informational robustness in intertemporal pricing." *Review of Economic Studies*, 88(3): 1224–1252.

Macchiavello, Rocco, and Ameet Morjaria. 2023. "Relational contracts: Recent empirical advancements and open questions." *Journal of Institutional and Theoretical Economics*, 179(3-4): 1–28.

Roesler, Anne-Katrin, and Balázs Szentes. 2017. "Buyer-optimal learning and monopoly pricing." *American Economic Review*, 107(7): 2072–2080.

Sobel, Joel, and Ichiro Takahashi. 1983. "A multistage model of bargaining." *Review of Economic Studies*, 50(3): 411–426.

Stokey, Nancy L. 1979. "Intertemporal price discrimination." *The Quarterly Journal of Economics*, 93(3): 355–371.

Thisse, Jacques-Francois, and Xavier Vives. 1988. "On the strategic choice of spatial price policy." *American Economic Review*, 78(1): 122–137.

Appendix B: Supplementary Results

Lemma B.1. *In equilibrium, both H-type buyers accept the first-period offer with the same probability σ.*

Proof. In what follows, we will make use of the following expressions for the posterior beliefs following good and bad news, extended to the case of asymmetric trading probabilities:

$$\pi_2(z_g;\sigma_i,\sigma_j) = \frac{\lambda[\sigma_i(1-\sigma_j)+(1-\sigma_i)\sigma_j]}{\lambda[\sigma_i(1-\sigma_j)+(1-\sigma_i)\sigma_j]+(1-\lambda)(\sigma_i+\sigma_j)}, \tag{B.1}$$

$$\pi_2(z_b;\sigma_i,\sigma_j) = \frac{\pi_1\lambda(1-\sigma_i)(1-\sigma_j)+\frac{\pi_1(1-\lambda)}{2}\left[(1-\sigma_i)+(1-\sigma_j)\right]}{\pi_1\lambda(1-\sigma_i)(1-\sigma_j)+\pi_1(1-\lambda)\left[(1-\sigma_i)+(1-\sigma_j)\right]+1-2\pi_1+\pi_1\lambda}. \tag{B.2}$$

We begin the proof by supposing, for contradiction, that $\sigma_i > \sigma_j$. This implies $\sigma_i > 0$ and $\sigma_j < 1$. We consider four cases, depending on the region to which the equilibrium first-period price belongs.

Case 1. Suppose that the seller posts a first-period offer in the following interval:

$$p_1 \in (v_H - \delta(v_H - v_L), v_H)\,.$$

Then, second-period prices must be contingent on news, so that second-period surpluses satisfy

$$S_H(z_g) < S_H(z_b).$$

To see why, note that if $S_H(z_g) = S_H(z_b) = 0$, then H-types are better off by trading for sure in the first period, contradicting $\sigma_j < 1$; if $S_H(z_g) = S_H(z_b) = v_H - v_L$, then H-types are better off delaying trade to the second period, contradicting $\sigma_i > 0$. And it is not possible to have $S_H(z_g) = S_H(z_b) \in (0, v_H - v_L)$ since $\sigma_i > 0$ implies that the posterior following good news is strictly higher than that following bad news ((B.1) and (B.2)). The same ordering gives $\phi(z_g) \geq \phi(z_b)$ and hence excludes $S_H(z_g) > S_H(z_b)$. The resulting continuation values of high-valuation buyers i and j are, respectively:

$$Q_H^i = \delta\left[\lambda\sigma_j S_H(z_g) + (1-\lambda\sigma_j)S_H(z_b)\right],$$
$$Q_H^j = \delta\left[\lambda\sigma_i S_H(z_g) + (1-\lambda\sigma_i)S_H(z_b)\right].$$

Since $\sigma_i > \sigma_j$ and $S_H(z_g) < S_H(z_b)$, it follows that $Q_H^i > Q_H^j$. However, $\sigma_i > 0$ requires $v_H - p_1 \geq Q_H^i$, whereas $\sigma_j < 1$ requires $v_H - p_1 \leq Q_H^j$, and these inequalities contradict $Q_H^i > Q_H^j$.

Case 2. Suppose that the seller posts the following first-period offer:

$$p_1 = v_H - \delta(v_H - v_L).$$

Since $\sigma_j < 1$, the seller can lower the offer slightly. At any price strictly below $v_H - \delta(v_H - v_L)$, trading immediately is strictly optimal for both high-valuation buyers, regardless of continuation play. Hence, both trading probabilities jump to one, and the discrete increase in first-period trade dominates the arbitrarily small reduction in the offer, contradicting seller optimality.

Case 3. It is also straightforward to rule out first-period offers satisfying

$$p_1 < v_H - \delta(v_H - v_L) \qquad \text{or} \qquad p_1 > v_H.$$

In the former case, the seller can profitably deviate and raise the price slightly without reducing the high type's probability of trade (which must equal 1), contradicting seller optimality. In the latter, buyer i would never accept any offer above v_H, contradicting buyer optimality.

Case 4. Suppose the seller posts a first-period offer $p_1 = v_H$. A high-valuation buyer who trades with positive probability must expect zero continuation surplus. Hence, the seller must offer v_H with probability one following bad news so that: $\pi_2(z_b; \sigma_i, \sigma_j) \geq \pi^\star$. As in the main text, let $\underline{\sigma}$ denote the symmetric trading probability satisfying $\pi_2(z_b; \underline{\sigma}, \underline{\sigma}) = \pi^\star$. Since, holding the sum $\sigma_i + \sigma_j$ fixed, the posterior $\pi_2(z_b; \sigma_i, \sigma_j)$ as defined in (B.2) is increasing in $(1 - \sigma_i)(1 - \sigma_j)$, we have that $\sigma_i \neq \sigma_j$ implies $\frac{1}{2}(\sigma_i + \sigma_j) < \underline{\sigma}$. Thus, under the asymmetric candidate equilibrium, the seller's revenue per buyer is

$$\pi_1 \left[\frac{\sigma_i + \sigma_j}{2} + \left(1 - \frac{\sigma_i + \sigma_j}{2} \right) \delta \right] v_H.$$

Instead, if the seller were to reduce the offer by a small amount $\varepsilon > 0$ below v_H, the trading probabilities would jump to $\underline{\sigma}$ and his revenue per buyer would be[1]

$$\pi_1 \left[\underline{\sigma} + (1 - \underline{\sigma}) \, \delta \right] v_H - \pi_1 \underline{\sigma} \varepsilon.$$

Since the increase in trading probabilities is discrete, the deviation is profitable for ε small enough, contradicting seller optimality. □

B.1 Continuous Valuations

We assume that each buyer's valuation v is continuously distributed on $[v_L, v_H]$, with $0 < v_L < v_H$. Conditional on an aggregate state $\omega \in \{h, \ell\}$, valuations are drawn from one of two distributions. More specifically, with probability π_1, both buyers' valuations are drawn from a distribution $F_h(\cdot)$ that is skewed toward high valuations; with probability $1 - \pi_1$, both are drawn from a distribution $F_\ell(\cdot)$ that is skewed toward low valuations. We consider the following parametric specification:

$$F_\omega(v) = \begin{cases} \left(\frac{v - v_L}{v_H - v_L} \right)^{\frac{1-\lambda}{1-\pi_1}} & \text{if } \omega = \ell, \\ \left(\frac{v - v_L}{v_H - v_L} \right)^{\frac{1-\pi_1}{1-\lambda}} & \text{if } \omega = h. \end{cases}$$

The above distributions admit densities $f_h(\cdot)$ and $f_\ell(\cdot)$, respectively. As in the baseline model, the parameter $\lambda \in [\pi_1, 1)$ governs the degree of correlation in valuations. When $\lambda = \pi_1$, the two conditional distributions coincide, so the aggregate state is uninformative and buyers' valuations are independent. As

[1] That trading probabilities are symmetric follows by the same argument as in *Case 1*. That they are equal to $\underline{\sigma}$ follows from the construction of the demand schedule in Section 3.1, where in the second period the seller offers v_H after good news with probability 1 and after bad news with some probability $\phi(z_b) \in (0, 1)$.

$\lambda \to 1$, the high-state distribution becomes concentrated at v_H while the low-state distribution becomes concentrated at v_L, so valuations become nearly perfectly correlated.

This specification preserves the information structure of the baseline model while replacing the binary valuation space with a continuum of valuations. In particular, observing one buyer's trading decision remains informative about the likely valuation of the other buyer, and the informativeness of this signal increases with λ. The key mechanism is therefore unchanged. A buyer's incentive to delay trade depends on the information that other buyers' trading decisions reveal about future demand. Consequently, the continuation value of waiting depends on the trading behavior of the other buyer, creating strategic complementarities similar to those in the baseline model.

We focus on symmetric equilibria in which buyers use cutoff strategies: given a first-period price p_1, a buyer trades immediately if and only if her valuation exceeds a cutoff $v^\star(p_1)$, at which she is indifferent between trading immediately and delaying trade. We later verify that such strategies are optimal. As in the baseline model, let $z \in \{z_g, z_b\}$ with $z_g = (g, p_1)$ and $z_b = (b, p_1)$ denote the news received by a buyer who delays, depending on whether the other buyer trades in the first period. The equilibrium probability of news z for a buyer with valuation v is given by:

$$\rho_v(z) = \gamma_v \rho_h(z) + (1 - \gamma_v)\rho_\ell(z),$$

where $\gamma_v = \frac{f_h(v)\pi_1}{f_h(v)\pi_1 + f_\ell(v)(1-\pi_1)}$ is the buyer's prior belief that the aggregate state is h and where:

$$\rho_\omega(z) = \begin{cases} 1 - F_\omega\left(v^\star(p_1)\right) & \text{if } z = z_g, \\ F_\omega\left(v^\star(p_1)\right) & \text{if } z = z_b, \end{cases}$$

is the probability of news z conditional on aggregate state ω. Thus, as in the baseline model, the distribution of news is endogenous to the seller's first-period price through its effect on the marginal type $v^\star$ who is just willing to trade in that period.

Conditional on history z, the seller updates his beliefs about the valuation of a buyer who delayed trade and chooses a second-period monopoly price:

$$p_2(z) \in \arg \max_{p \in [v_L,\, v_H]} p\left[1 - F_2(p; z)\right],$$

where $F_2(\cdot; z)$ is the distribution that the seller faces in the second period, given history z:

$$F_2(v; z) = \begin{cases} \frac{F_h(v)\rho_h(z)\pi_1 + F_\ell(v)\rho_\ell(z)(1-\pi_1)}{F_h(v^\star)\rho_h(z)\pi_1 + F_\ell(v^\star)\rho_\ell(z)(1-\pi_1)} & \text{if } v < v^\star(p_1), \\ 1 & \text{if } v \geq v^\star(p_1). \end{cases}$$

The resulting v-type buyer's surplus in the second period, after trading history z, can thus be expressed as

$$S_v(z) = \begin{cases} 0 & \text{if } v \leq p_2(z), \\ v - p_2(z) & \text{if } v > p_2(z). \end{cases}$$

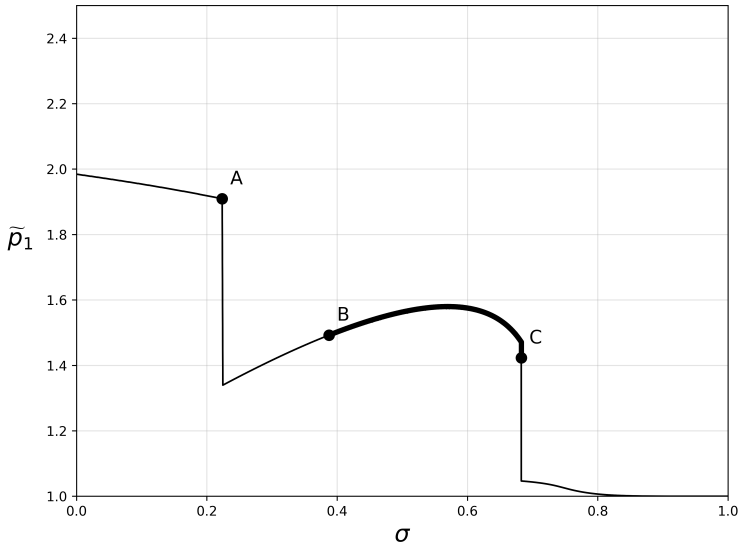


Figure B.1: **The inverse demand schedule with continuous valuations**. In this figure, the parameters used are: $\pi_1 = 0.7$, $v_L = 1$, $v_H = 2.5$, $\delta = 0.95$, and $\lambda = 0.95$.

Hence, the cutoff buyer who is indifferent about accepting a first-period offer p_1 is given by:

$$v^\star(p_1) - p_1 = \delta \left[\rho_{v^\star(p_1)}(z_g) S_{v^\star(p_1)}(z_g) + (1 - \rho_{v^\star(p_1)}(z_g)) S_{v^\star(p_1)}(z_b)\right]. \tag{B.3}$$

The analogue of the skimming property in Lemma 1 also holds here. Holding fixed the continuation play induced by a candidate cutoff, the left-hand side of (B.3), evaluated at any valuation v, has slope one in v, whereas the right-hand side has slope at most $\delta < 1$. To see this, note that, for $\lambda > \pi_1$, f_h/f_ℓ is increasing and $F_h < F_\ell$, so $\rho_v(z_g)$ increases in v. The seller's posterior after good news dominates that after bad news in the likelihood-ratio order, so $p_2(z_g) \geq p_2(z_b)$ and hence $S_v(z_g) \leq S_v(z_b)$. Finally, each $S_v(z)$ has slope at most one in v. Thus, given any $v^\star(p_1)$ satisfying (B.3), buyers with higher valuations strictly prefer to trade immediately and those with lower valuations prefer to delay.

Equation (B.3) also defines the inverse demand schedule faced by the seller in the first period. As in the baseline model, the shape of this schedule depends on how buyers' trading decisions affect the information available to the seller in the second period. Figure B.1 plots the inverse demand schedule implied by (B.3). For each first-period price $\widetilde{p}_1$, the corresponding value $\sigma(\widetilde{p}_1)$ denotes the equilibrium probability that a buyer's valuation exceeds the cutoff $v^\star(\widetilde{p}_1)$.

As in the baseline model, sufficiently strong correlation generates an upward-sloping region of the inverse demand schedule. The intuition is identical. When buyers expect more first-period trade, observing another buyer purchase becomes more likely. This makes the seller more optimistic about the valuation of buyers who delay, inducing a higher continuation price. Since waiting becomes less attractive, buyers become more willing to purchase immediately. Consequently, buyers' trading decisions are strategic complements. These complementarities imply that demand is nonmonotone. Over the upward-sloping region, an increase in the expected probability of first-period trade raises buyers' willingness to purchase immediately, so a higher first-period price can be associated with greater demand.

Finally, these strategic complementarities also give rise to equilibrium multiplicity. Verifying equilibrium analytically is more involved than in the binary model because the seller faces an interior demand elasticity throughout. Nevertheless, we can compute the equilibrium set numerically. For the parameterization in Figure B.1, point A and every point on the thick segment B-C of the inverse demand schedule are equilibria: given buyers' responses to an off-equilibrium price,[2] the seller has no profitable deviation. Hence, the same mechanism as in the baseline model generates strategic complementarities, nonmonotone demand, and equilibrium multiplicity with continuous valuations.

B.2 Many Buyers

We now extend the baseline model to an arbitrary number of buyers. Suppose there are $N+1$ buyers whose valuations are correlated through a common aggregate state $\omega \in \{h, \ell\}$. As in the baseline model, the unconditional distribution of types is

$$\mathbb{P}(\theta_i = H) = \pi_1.$$

The aggregate state satisfies

$$\mathbb{P}(\omega = h) = \pi_1,$$

and, conditional on the aggregate state, for every buyer i,

$$\mathbb{P}(\theta_i = H|\omega = h) = \lambda, \qquad \mathbb{P}(\theta_i = H|\omega = \ell) = \frac{(1-\lambda)\pi_1}{1-\pi_1}.$$

Thus, types are independent at $\lambda = \pi_1$, and correlation increases in λ, becoming perfect as $\lambda \to 1$.[3]

The analysis closely parallels that of the baseline model. Consider a buyer i who delays trade to the second period. From her perspective, the relevant news is now how many of the remaining N buyers trade in the first period. Let z_n denote the event that exactly $n \in \{0, \ldots, N\}$ of these buyers trade. If in equilibrium a high-valuation buyer accepts the first-period offer with probability σ, then conditional on buyer i's type, the probability of observing the news realization z_n is

$$\rho_\theta(z_n; \sigma) = \sum_{\widetilde{\omega} \in \{\ell, h\}} \mathbb{P}(\omega = \widetilde{\omega}|\theta_i = \theta) \binom{N}{n} q(\widetilde{\omega})^n \big(1 - q(\widetilde{\omega})\big)^{N-n},$$

where

$$q(\widetilde{\omega}) = \sigma\, \mathbb{P}(\theta_j = H|\omega = \widetilde{\omega})$$

is the probability that any one of the remaining buyers trades in the first period conditional on the aggregate state. As in the baseline model, the news distribution is endogenous because it depends on the trading probability σ, which in turn is determined by the seller's first-period price.

Following a history z_n, the seller updates his posterior belief $\pi_2(z_n; \sigma)$ that buyer i has a high valuation using Bayes' rule. As before, he charges v_H whenever $\pi_2(z_n; \sigma) > \pi^\star$, charges v_L whenever $\pi_2(z_n; \sigma) < \pi^\star$,

[2] Following an off-path price, buyers coordinate on the lowest trading probability consistent with (B.3).

[3] In this extension, λ is the probability of a high type in state h. The implied correlation between any two buyers is $\mathbb{P}(\theta_j = H|\theta_i = H) = \pi_1 + (\lambda - \pi_1)^2/(1-\pi_1)$ and coincides with λ only at $\lambda = \pi_1$ and as $\lambda \to 1$.

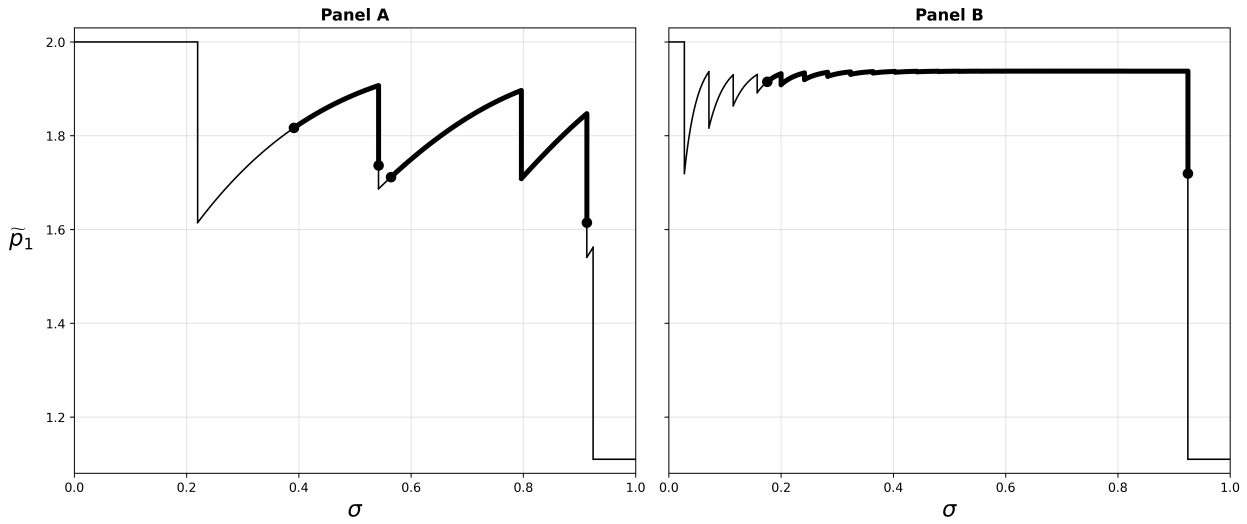


Figure B.2: **Inverse demands with many buyers.** The inverse demand schedule and contingent-price equilibrium set (thick segments) with $N+1=5$ buyers (Panel A) in comparison to the inverse demand schedule with $N+1=50$ buyers (Panel B). In this figure, $\pi_1=0.7$, $v_L=1$, $v_H=2$, $\lambda=0.93$, $\delta=0.89$.

and mixes between the two prices with probability $\phi(z_n)$ whenever $\pi_2(z_n;\sigma)=\pi^\star$. Since buyers' continuation values remain given by (7) and (8), the inverse demand schedule satisfies:

$$v_H - p_1 = \delta \sum_{n=0}^{N} \rho_H(z_n;\sigma)\big(1-\phi(z_n)\big)(v_H - v_L).$$

The key difference relative to the baseline is that the seller's first-period price now affects buyers' continuation values through a richer distribution of future news realizations. By changing the first-period trading probability σ, a higher price changes the likelihood of each news realization, the seller's posterior beliefs, and hence his continuation pricing strategy. This feedback between current prices, future beliefs, and buyers' continuation values gives rise to the same strategic complementarities in first-period purchase decisions: a buyer becomes more willing to trade early when she expects other buyers to do so as well. Consequently, the seller may again face a nonmonotone demand schedule, and the same mechanism generates equilibrium multiplicity.

Figure B.2 illustrates the inverse demand schedule for a market with $N+1=5$ buyers (Panel A) and a market with $N+1=50$ buyers (Panel B). Although the richer news space makes a complete analytical characterization of the equilibrium set considerably more involved, it is straightforward to compute the equilibria numerically. For the parameterization in Figure B.2, the equilibrium set lies within the thick segments.[4] Thus, the same strategic complementarities as in the baseline model continue to generate nonmonotone demand and equilibrium multiplicity in large markets.

[4] Although the thick shaded segment in Panel B appears connected, not every point on it is in fact an equilibrium; this is difficult to discern visually.

B.3 Common Values

In the baseline model, a buyer's valuation depends only on her own type θ_i. Trading histories therefore reveal information about the other buyer's valuation, but they are of no direct payoff relevance for the buyer's own consumption value. Here we introduce a common-value component to valuations and argue that our results are robust to a small perturbation in this direction.

Suppose each buyer i privately observes a signal $\theta_i \in \{H, L\}$ at $t = 1$, distributed exactly as in Section 2. The buyer learns her true valuation only after trade takes place. With probability $1 - \gamma$ it equals v_{θ_i}, and with probability $\gamma \in [0, \frac{1}{2})$ it instead equals v_{θ_j}. Thus, θ_i remains the buyer's private information at the trading stage, but her valuation is tied to the other buyer's signal with probability γ. When $\gamma = 0$, this specification reduces to our baseline model. For $\gamma > 0$, buyer i's expected valuation conditional on her own signal is:

$$\mathbb{E}[v_i|\theta_i] = (1 - \gamma)v_{\theta_i} + \gamma\mathbb{E}[v_{\theta_j}|\theta_i], \tag{B.4}$$

where:

$$\mathbb{E}[v_{\theta_j}|\theta_i] = \mathbb{P}(\theta_j = H|\theta_i)v_H + \big(1 - \mathbb{P}(\theta_j = H|\theta_i)\big)v_L. \tag{B.5}$$

Similarly, a buyer who delays trade and observes first-period history z values the good at $\mathbb{E}[v_i|\theta_i, z]$, where the first-period history z now also affects her beliefs about θ_j.

The buyers' expected valuations, the seller's posterior beliefs, and hence the second-period pricing problem and the buyers' indifference conditions that define the inverse demand schedule all vary continuously in γ. Thus, our qualitative results continue to hold when γ is small: whenever the baseline model features strategic complementarities, an upward-sloping region of inverse demand, and multiple equilibria, with the relevant inequalities holding strictly (e.g., $\lambda > \lambda^{nc}$), so does the model with a sufficiently small common-value component.

When γ is not small, however, new forces emerge that are absent from our baseline analysis. Most importantly, trading histories now affect not only the seller's beliefs about buyers' willingness to pay, but also buyers' own expected valuations. If a buyer rejects the first-period offer, she is more likely to have a low signal. Since signals now partially determine both buyers' realized valuations, the other buyer's expected valuation also declines. As a result, second-period reservation values become more sensitive to first-period rejections than in our baseline model. This induces the seller to lower the initial price more aggressively in order to avoid such rejections. This additional informational channel is absent under correlated private values, and it interacts with the seller's dynamic pricing incentives in a nontrivial way. We conjecture that our central mechanism continues to operate in this richer environment, but characterizing the interaction between these two informational channels is beyond the scope of the present paper.

B.4 Low-Prior Case

In our baseline model, we imposed Assumption 1, which ensures that pooling allocations do not arise in the corresponding static pricing problem. We now argue that the mechanism generating strategic complementarities, nonmonotone demand, and equilibrium multiplicity does not rely on this assumption. In

particular, it also operates in the low-prior case, $\pi_1 v_H < v_L$, provided that the seller is sufficiently patient. Since the seller's revenue in any contingent-price equilibrium is bounded above by $\pi_1 v_H + (1-\pi_1)\delta v_L$, these equilibria can exist only if this bound exceeds the pooling revenue v_L; we therefore restrict attention to $\delta > (v_L - \pi_1 v_H)/[(1-\pi_1)v_L]$.

In this case, when correlation is sufficiently weak, i.e., $\lambda < \widehat{\lambda}^{nc}$, the seller optimally trades with both buyer types immediately at price v_L.[5] This contrasts with the baseline setting, in which either a high-price equilibrium (where only high-value buyers trade) or a low-price equilibrium (where both types trade, but with delay) arises. Once correlation exceeds the threshold, $\lambda > \widehat{\lambda}^{nc}$, the same mechanism as in the baseline model makes the inverse demand schedule nonmonotone and generates a continuum of contingent-price equilibria, in which second-period prices depend on the first-period trading history. Moreover, these equilibria coexist with the pooling equilibrium in which all buyers trade immediately at price v_L.

Figure B.3 illustrates the inverse demand schedule in the low-prior case, both without information spillovers ($\lambda = \pi_1$) and with information spillovers ($\lambda > \pi_1$). Without spillovers, the second-period price is always v_L, so no buyer trades in the first period at price $p_1 > v_H - \delta(v_H - v_L)$. At the cutoff price $p_1 = v_H - \delta(v_H - v_L)$, the high-value buyer is indifferent, so any $\sigma_H \in [0,1]$ is consistent with equilibrium, where σ_H represents the high type's trading probability. For lower prices, the high-value buyer trades with a probability of one, while the low-valuation buyer also trades at a price $p_1 \leq v_L$, i.e., $\sigma_L \geq 0$, with a probability $\sigma_L = 1$ if $p_1 < v_L$.

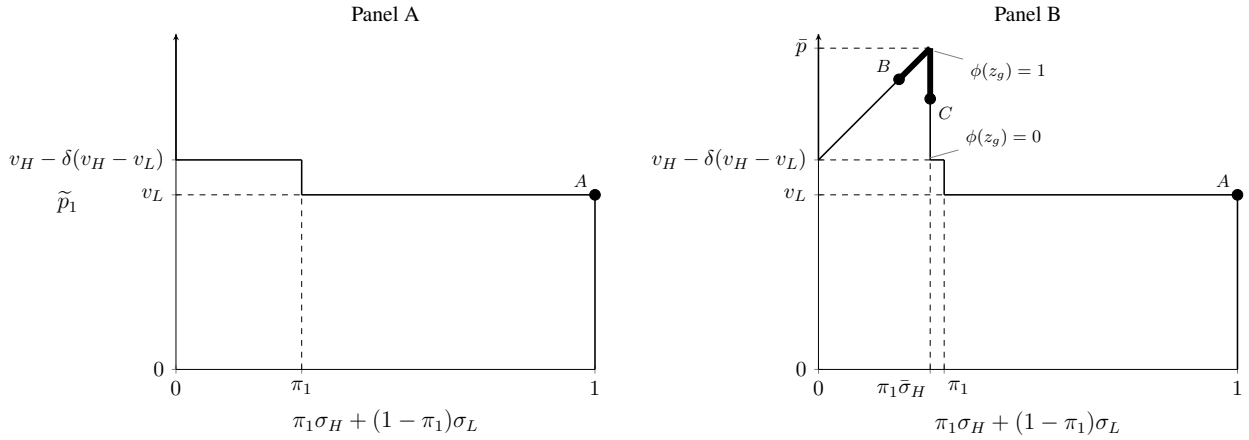


Figure B.3: **The inverse demand schedule in the low-prior case**. Panel A depicts the case of no information spillovers, whereas Panel B depicts the case with information spillovers. In both panels, we set $\pi_1 = 0.3$, $v_L = 1$, $v_H = 2$ and $\delta = 0.8$. Moreover, in Panel A we set $\lambda = 0.3$ and in Panel B we set $\lambda = 0.9$.

When information spillovers are sufficiently strong, however, a first-period trade may induce the seller to charge a second-period price above v_L. Anticipation of these higher future prices strengthens buyers' incentives to trade early, giving rise to nonmonotone demand, equilibrium multiplicity, and the ability of the seller to sustain higher first-period prices. For the parameterization in Figure B.3, the equilibrium set consists of every point on the thick shaded segment of the inverse demand schedule (segment B–C), together with the pooling equilibrium in which all buyers trade immediately at price v_L (point A).

[5] In the low-prior case, $\widehat{\lambda}^{nc} \in (\pi^\star, 1)$ is the unique value of λ solving $R^{cp}(\bar{p}, \bar{\sigma}; \lambda) = v_L$.

B.5 Infinite Horizon

Our baseline model assumes two trading opportunities. We now consider an infinite-horizon extension in which the seller can post a new price in every period until trade occurs. The only substantive difference relative to the baseline environment is that there is no final trading period. Hence, the discount factor δ now also governs the frequency of trading opportunities and, therefore, the severity of the seller's commitment problem. As δ increases, trading opportunities become more frequent, and the seller becomes less able to sustain high future prices.

Absent information spillovers, our environment reduces to the classic dynamic monopoly problem studied by Fudenberg, Levine and Tirole (1985) and Gul, Sonnenschein and Wilson (1986). Because the seller cannot commit to future prices, buyers have an incentive to delay trade in anticipation of future price cuts, while the seller lowers prices over time to induce additional purchases. This dynamic commitment problem depresses the seller's revenues relative to commitment. As trading opportunities become arbitrarily frequent ($\delta \to 1$), these incentives culminate in the Coase conjecture: equilibrium prices and seller revenues converge to v_L.

When information spillovers are present, however, continuation play becomes substantially richer. If one buyer has already purchased, the seller faces the standard one-buyer screening problem. If instead both buyers remain active, current trade affects the seller's beliefs about the remaining buyer and therefore the entire continuation price path. This informational feedback gives rise to the same strategic complementarities as in the baseline model.

We now show that, as in our baseline setting, sufficiently strong information spillovers undermine the Coasian forces. Let $R^c(1) \equiv \pi_1 v_H + (1-\pi_1)\delta v_L$ denote the seller's commitment revenue per buyer as $\lambda \to 1$.

Proposition B.1. *Fix $\delta < 1$. As $\lambda \to 1$, there exists a sequence of equilibria indexed by λ along which: (i) high-type buyers trade in the first period at a price converging to v_H and with probability converging to one; (ii) low-type buyers trade in the second period at price v_L with probability converging to one; and, consequently, (iii) the seller's expected discounted revenues converge to $R^c(1)$.*

The proof (see below) constructs an equilibrium in which high-type buyers trade almost immediately when valuations are sufficiently correlated. The key observation is that, following a purchase by one buyer, the seller faces a standard one-buyer bargaining problem with a posterior that becomes arbitrarily close to one as $\lambda \to 1$. The remaining high-type buyer then obtains arbitrarily little continuation surplus. By contrast, following no trade, the seller's posterior remains sufficiently pessimistic that he charges v_L in the next period. A high-type buyer who delays therefore faces an increasingly likely transition to a continuation game in which the seller extracts almost all of her surplus. This allows the seller to charge an initial price arbitrarily close to v_H while inducing high types to trade with probability arbitrarily close to one.

This construction also illustrates that information spillovers generate the same strategic complementarity as in the baseline model. When other high-type buyers trade more aggressively, an observed purchase becomes more informative, raising the seller's posterior about any remaining buyer and strengthening his incentives to charge high continuation prices. The resulting decline in the payoff from waiting, in turn, makes each high-type buyer more willing to trade early.

As in the baseline model, however, strong spillovers do not guarantee that the seller approaches his commitment payoff in all equilibria.[6]

Proposition B.2. *Fix $\delta < 1$. As $\lambda \to 1$, there also exists a sequence of equilibria indexed by λ along which the seller's expected discounted revenues remain bounded away from $R^c(1)$.*

The proof (also below) constructs an equilibrium in which high types trade in the first period with a probability $\sigma_1 < 1$ that does not depend on λ. A purchase still makes the seller confident that the remaining buyer is a high type, but with probability bounded away from zero no purchase occurs, in which case the seller pools at v_L. Waiting thus remains valuable for high types, which keeps the first-period price, and hence the seller's revenue, bounded away from the commitment benchmark.

If the seller deviates to an offer that both buyers reject for sure, nothing is learned, so the game restarts one period later and we let this same equilibrium be played again. Since high types keep a payoff bounded away from zero, rejecting offers close to v_H is optimal. After other offers, we specify continuation play using the information spillovers (Lemmas B.2–B.4). A deviation followed by certain rejection only delays the equilibrium revenue. Any other deviation either delays trade with high types or leaves them a rent, and so yields at most a fixed amount strictly below the commitment benchmark. When σ_1 is close enough to one, the equilibrium revenue exceeds this amount, and no deviation is profitable.

Preliminaries. Let $V(\pi)$ denote the seller's expected discounted revenue in the equilibrium of the one-buyer continuation game starting from belief π, and let $U(\pi)$ denote the corresponding payoff of the high-type buyer. Standard results from the dynamic bargaining literature (e.g.; Fudenberg, Levine and Tirole, 1985; Gul, Sonnenschein and Wilson, 1986) imply that $V(\pi) = v_L$ and $U(\pi) = v_H - v_L$ whenever $\pi < \pi^\star$, while $V(\pi) \to v_H$ and $U(\pi) \to 0$ as $\pi \to 1$. Moreover, no price below v_L is ever posted, and a low type never pays more than v_L. At the threshold beliefs of the one-buyer game the seller may randomize over his first offer, so U is a correspondence with values in $[0, v_H - v_L]$ whose graph is connected, while V is continuous. We write U_σ for a value of $U(\pi_2(z_g; \sigma))$.

If high types accept a first-period offer with probability σ and low types reject it, the posterior about a remaining buyer is $\pi_2(z_g; \sigma) = \lambda(1-\sigma)/(1-\lambda\sigma)$ after good news and $\pi_2(z_b; \sigma)$ after bad news, as in the baseline model. As in the baseline, $\underline{\sigma} \in (0,1)$ solves $\pi_2(z_b; \underline{\sigma}) = \pi^\star$. It decreases in λ and converges to $\underline{\sigma}(1) \in (0,1)$ as $\lambda \to 1$. Hence $\pi_2(z_g; \underline{\sigma}) \to 1$. Fix any $U_{\underline{\sigma}} \in U(\pi_2(z_g; \underline{\sigma}))$, so that $U_{\underline{\sigma}} \to 0$, and define $Q_{\underline{\sigma}} \equiv \lambda\underline{\sigma}\, U_{\underline{\sigma}} + (1 - \lambda\underline{\sigma})(v_H - v_L)$. Note that $Q_{\underline{\sigma}} \to (1 - \underline{\sigma}(1))(v_H - v_L)$ as $\lambda \to 1$. Finally, for any punishment revenue per buyer R_D, let

$$\mathcal{D}(\lambda; R_D) \equiv \max\left\{\delta R_D,\ R^c(1) - \pi_1(1-\delta)(1-\underline{\sigma})v_H,\ R^c(1) - \delta\pi_1 Q_{\underline{\sigma}},\ v_L\right\}. \tag{B.6}$$

Lemma B.2. *Consider any history at which both buyers remain and the seller's belief about each is some $\pi \le \pi^\star$, whatever the correlation between them.*

[6] Propositions B.1 and B.2 show that, as $\lambda \to 1$, the seller's revenue converges to the commitment revenue in some equilibria, whereas it remains bounded away from it in others. Note that, in the equilibria of Proposition B.2, some high types delay trade, so social surplus is also lower than in those of Proposition B.1. A characterization and analysis of the entire equilibrium set is considerably more involved in the infinite-horizon environment. Following a given initial offer, there are now histories at which both buyers remain active and information spillovers persist, but the seller's beliefs have drifted downward after successive rejections.

(i) The following pooling continuation is an equilibrium: the seller offers v_L in every period; after any offer $\tilde{p}$, high types accept if and only if $\tilde{p} \leq v_H - \delta(v_H - v_L)$, and low types if and only if $\tilde{p} \leq v_L$. In it, no offer yields the seller more than v_L.

(ii) Any other continuation from such a history, in which the seller first posts some prescribed offer, is also an equilibrium if: (a) buyers best respond to that offer; (b) after it, play from the next period on is a continuation equilibrium; (c) after any other offer, play follows the pooling continuation in (i); and (d) the seller earns at least v_L.

Proof. Since no later price is below v_L, the stated responses are optimal. After an offer that high types accept for sure, any remaining buyer is known to be a low type and is offered v_L. After an offer that both types reject, the belief is unchanged. Hence the specification is self-consistent. An offer $\tilde{p} \geq v_H - \delta(v_H - v_L)$ is rejected and yields δv_L. An offer $\tilde{p} \in \big(v_L, v_H - \delta(v_H - v_L)\big)$ yields $\pi\tilde{p} + (1-\pi)\delta v_L < \pi(1-\delta)v_H + \delta v_L \leq v_L$, where the last inequality uses $\pi \leq \pi^\star = v_L/v_H$. An offer $\tilde{p} \leq v_L$ yields at most v_L. For (ii), any deviation by the seller is followed by pooling, so by (i) it yields at most v_L, which is no more than he earns by following the continuation. □

Lemma B.3. *Suppose high types accept a first-period offer with probability $\underline{\sigma}$, low types reject it, and neither buyer trades. For λ close enough to one, there is $\underline{Q}_b(\lambda)$, with $\underline{Q}_b(\lambda) \to 0$ as $\lambda \to 1$, such that for every $Q_b \in [\underline{Q}_b(\lambda), v_H - v_L]$ the continuation game has an equilibrium in which the high type's payoff, at the start of the second period, is Q_b.*

Proof. The seller's belief about each remaining buyer is $\pi^\star$, and the probability that one is a high type given that the other is equals $\lambda' \equiv \lambda(1-\underline{\sigma})/(1-\lambda\underline{\sigma})$, which converges to one.

For any $\sigma_2 \in (0,1]$, consider the following continuation. The seller offers a price that high types accept with probability σ_2 and low types reject. After one purchase, the one-buyer game is played from belief $\pi_3(\sigma_2) \equiv \lambda'(1-\sigma_2)/(1-\lambda'\sigma_2)$. Since $\lambda' > \pi^\star$, for λ close to one, the belief after no purchase is below $\pi^\star$ and the seller pools as in Lemma B.2(i). After any other second-period offer, play follows Lemma B.2(i). High types are indifferent at the price $p_2 = v_H - Q_b$, where $Q_b = \delta[\lambda'\sigma_2 U' + (1-\lambda'\sigma_2)(v_H - v_L)]$ and $U' \in U(\pi_3(\sigma_2))$. Since $Q_b \leq \delta(v_H - v_L)$, $p_2 > v_L$ and low types reject. The seller's revenue R_2 and the high type's payoff Q_b are continuous on the graph $\{(\sigma_2, U') : U' \in U(\pi_3(\sigma_2))\}$, which is connected. By Lemma B.2(ii), each such profile is a continuation equilibrium whenever $R_2 \geq v_L$.

Let $1 - \sigma_2 = \sqrt{1-\lambda'}$. As $\lambda \to 1$, $\pi_3(\sigma_2) \to 1$ and $1 - \lambda'\sigma_2 \to 0$, so $Q_b \to 0$. A high type trades with probability converging to one at a price converging to v_H. A low type faces a purchase by the other buyer with probability $\pi^\star(1-\lambda')\sigma_2/(1-\pi^\star) \to 0$, and otherwise pays v_L in the next period. Hence $R_2 \to \pi^\star v_H + (1-\pi^\star)\delta v_L = v_L + (1-\pi^\star)\delta v_L > v_L$. For λ close to one, let $\underline{Q}_b(\lambda)$ be the high type's payoff at this point, so $\underline{Q}_b(\lambda) \to 0$.

As $\sigma_2 \to 0$, $R_2 \to \delta v_L < v_L$. Moving along the graph from the point with $1 - \sigma_2 = \sqrt{1-\lambda'}$ toward $\sigma_2 = 0$, consider the first point at which $R_2 = v_L$, and let Q_b° be the high type's payoff there. Up to that point $R_2 \geq v_L$, so every profile is a continuation equilibrium, and by connectedness they deliver every Q_b between $\underline{Q}_b(\lambda)$ and Q_b°. At that point the seller earns v_L, as he does by pooling, so he is indifferent and may randomize between the two second-period offers. This delivers every Q_b between Q_b° and $v_H - v_L$. □

Lemma B.4. *Fix λ close enough to one for Lemma B.3 to apply, and let*

$$\underline{Q}(\lambda) \equiv \lambda \underline{\sigma}\, U_{\underline{\sigma}} + (1 - \lambda \underline{\sigma}) \underline{Q}_b(\lambda), \tag{B.7}$$

so that $\underline{Q}(\lambda) \to 0$. Suppose that, after any first-period offer that both types reject for sure, play continues from the next period with a profile that gives the high type a payoff $Q_D \geq \underline{Q}(\lambda)$ and the seller revenue R_D. Then, after every first-period offer $\tilde{p}$, there is a specification of buyers' responses and subsequent play in which buyers best respond, play is a continuation equilibrium at every later history except those following an offer that both types reject for sure, and the seller's revenue is at most $\mathcal{D}(\lambda; R_D)$.

Proof. Let $Q = v_H - \tilde{p}$ denote the high type's surplus from accepting. We assign each offer to one region.

Region 1: $Q \leq \delta Q_D$. Both types reject, which is optimal for a high type because waiting yields $\delta Q_D \geq Q$, and for a low type because $\tilde{p} \geq v_H - \delta(v_H - v_L) > v_L$. Nothing is learned, and the seller earns δR_D.

Region 2: $\delta Q_D < Q < \delta Q_{\underline{\sigma}}$. High types accept with probability $\underline{\sigma}$ and low types reject. After a purchase, the one-buyer game is played from $\pi_2(z_g; \underline{\sigma})$, with the selection $U_{\underline{\sigma}}$. After no purchase, play follows a continuation from Lemma B.3 with high-type payoff Q_b, chosen so that $Q = \delta[\lambda \underline{\sigma}\, U_{\underline{\sigma}} + (1 - \lambda \underline{\sigma}) Q_b]$. Such a $Q_b \in [\underline{Q}_b(\lambda), v_H - v_L]$ exists because Q lies between $\delta \underline{Q}(\lambda) \leq \delta Q_D$ and $\delta Q_{\underline{\sigma}}$. A fraction $1 - \underline{\sigma}$ of high types trades no earlier than the second period, so the seller earns at most $\pi_1[\underline{\sigma} v_H + (1 - \underline{\sigma})\delta v_H] + (1 - \pi_1)\delta v_L$, which equals $R^c(1) - \pi_1(1 - \delta)(1 - \underline{\sigma}) v_H$.

Region 3: $\delta Q_{\underline{\sigma}} \leq Q < \delta(v_H - v_L)$. High types accept with some probability $\sigma \in [\underline{\sigma}, 1]$ and low types reject. After a purchase, the one-buyer game is played from $\pi_2(z_g; \sigma)$. After no purchase, the belief is at most $\pi^\star$, and the seller pools (Lemma B.2). High types are indifferent if $Q = \delta[\lambda \sigma U_\sigma + (1 - \lambda\sigma)(v_H - v_L)]$ for some $U_\sigma \in U(\pi_2(z_g; \sigma))$. The right-hand side equals $\delta Q_{\underline{\sigma}}$ at $(\underline{\sigma}, U_{\underline{\sigma}})$ and $\delta(v_H - v_L)$ at $\sigma = 1$, where $\pi_2(z_g; 1) = 0$. Since it is continuous on a connected graph, it takes every value in between.

Region 4: $Q \geq \delta(v_H - v_L)$. High types accept for sure, since no later price is below v_L. Low types accept if and only if $\tilde{p} \leq v_L$. Any remaining buyer is known to be a low type and is offered v_L.

In regions 3 and 4 with $\tilde{p} > v_L$, low types trade no earlier than the second period, so total surplus is at most $R^c(1)$, while a high type obtains at least $Q \geq \delta Q_{\underline{\sigma}}$, since she can accept. Hence the seller earns at most $R^c(1) - \delta \pi_1 Q_{\underline{\sigma}}$. If $\tilde{p} \leq v_L$, he earns at most v_L. Later histories are continuation equilibria by Lemmas B.2 and B.3 and the equilibrium of the one-buyer game. □

Proof of Proposition B.2. Fix $\sigma_1 \in (\underline{\sigma}(1), 1)$, and let λ be close enough to one that $\sigma_1 > \underline{\sigma}$. Consider the following profile $\hat{e}$. At every history at which both buyers remain and every earlier offer was rejected for sure by both, beliefs equal the prior, and the seller offers $\hat{p} = v_H - \hat{Q}$, where $\hat{Q} \equiv \delta[\lambda \sigma_1 U_{\sigma_1} + (1 - \lambda\sigma_1)(v_H - v_L)]$ and $U_{\sigma_1} \in U(\pi_2(z_g; \sigma_1))$. High types accept with probability σ_1 and low types reject. After a purchase, the one-buyer game is played. After no purchase, the belief is at most $\pi^\star$ and the seller pools. After any other first-period offer, play follows Lemma B.4, where rejection for sure by both restarts $\hat{e}$, so that $Q_D = \hat{Q}$ and $R_D = \hat{R}$, the seller's revenue in $\hat{e}$.

Since $\pi_2(z_g; \sigma_1) \to 1$, $U_{\sigma_1} \to 0$ and $\hat{Q} \to \delta(1 - \sigma_1)(v_H - v_L) > 0$. A low type faces a purchase by the other buyer with probability converging to zero, and otherwise the seller pools and she pays v_L in the second

period. A high type who rejects faces a purchase by the other buyer with probability converging to σ_1, after which she pays approximately v_H, and otherwise pays v_L, both one period later. Hence

$$\hat{R} \to \pi_1\big[\sigma_1(v_H - \lim \hat{Q}) + (1-\sigma_1)\delta(\sigma_1 v_H + (1-\sigma_1)v_L)\big] + (1-\pi_1)\delta v_L = \pi_1\sigma_1 v_H + \delta(1-\pi_1\sigma_1)v_L,$$

which equals $R^c(1) - \pi_1(1-\sigma_1)(v_H - \delta v_L) < R^c(1)$. Also, $\mathcal{D}(\lambda; \hat{R})$ in (B.6) converges to

$$\max\big\{\delta \lim \hat{R},\ R^c(1) - \pi_1(1-\delta)(1-\underline{\sigma}(1))v_H,\ R^c(1) - \delta\pi_1(1-\underline{\sigma}(1))(v_H - v_L),\ v_L\big\}.$$

As $\sigma_1 \to 1$, $\lim \hat{R} \to R^c(1)$, which strictly exceeds each term, since $\underline{\sigma}(1) < 1$ and $v_L < \pi_1 v_H < R^c(1)$. Hence there is $\sigma_0 \in (\underline{\sigma}(1), 1)$ such that, for every $\sigma_1 \in [\sigma_0, 1)$, $\lim \hat{R}$ strictly exceeds the limit of $\mathcal{D}(\lambda; \hat{R})$.

Now fix such a σ_1, and let λ be close enough to one that $\hat{Q} \geq \underline{Q}(\lambda)$, with $\underline{Q}(\lambda)$ as in (B.7), and $\hat{R} \geq \mathcal{D}(\lambda; \hat{R})$. The first holds because $\hat{Q}$ stays bounded away from zero while $\underline{Q}(\lambda) \to 0$, and the second because the strict inequality holds in the limit. Buyers best respond on path by construction, and off path by Lemma B.4. At every history at which both buyers remain and every earlier offer was rejected for sure by both, the seller earns $\hat{R}$ on path and at most $\mathcal{D}(\lambda; \hat{R}) \leq \hat{R}$ from any other offer. At every other history play is a continuation equilibrium. Since $\delta < 1$ and, without loss, offers lie in $[0, v_H]$, payoffs are bounded and the one-shot deviation principle applies, so $\hat{e}$ is an equilibrium. Along these equilibria, the seller's revenue per buyer converges to a value strictly below $R^c(1)$. □

Proof of Proposition B.1. Let $\hat{e}$ be the equilibrium constructed in the proof of Proposition B.2. Take any sequence $\sigma_1(\lambda) \to 1$ that converges to one sufficiently slowly that $(1-\sigma_1(\lambda))/(1-\lambda) \to \infty$, for example $1 - \sigma_1(\lambda) = \sqrt{1-\lambda}$. Then $\sigma_1(\lambda) > \underline{\sigma}$ for λ close to one, and $\pi_2(z_g; \sigma_1(\lambda)) \to 1$ as well as $1 - \lambda\sigma_1(\lambda) \to 0$.

At period one, the seller offers $p_1(\lambda) \equiv v_H - \delta[\lambda\sigma_1(\lambda)U_{\sigma_1(\lambda)} + (1 - \lambda\sigma_1(\lambda))(v_H - v_L)]$, with $U_{\sigma_1(\lambda)} \in U(\pi_2(z_g; \sigma_1(\lambda)))$. High types accept with probability $\sigma_1(\lambda)$ and low types reject. After a purchase, the one-buyer game is played. After no purchase, the belief is at most $\pi^\star$ and the seller pools. After any other first-period offer, play follows Lemma B.4, where rejection for sure by both leads to the equilibrium $\hat{e}$ of Proposition B.2, so that $Q_D = \hat{Q}$ and $R_D = \hat{R}$.

Since $U_{\sigma_1(\lambda)} \to 0$ and $1 - \lambda\sigma_1(\lambda) \to 0$, $v_H - p_1(\lambda) \to 0$, which proves (i). A low type faces a purchase by the other buyer with probability $\pi_1(1-\lambda)\sigma_1(\lambda)/(1-\pi_1) \to 0$, and otherwise the seller pools and she pays v_L in the second period, which proves (ii). Hence the seller's revenue $R_1(\lambda)$ is at least

$$\pi_1\sigma_1(\lambda)p_1(\lambda) + (1-\pi_1)\delta v_L\Big[1 - \frac{\pi_1(1-\lambda)\sigma_1(\lambda)}{1-\pi_1}\Big] \to R^c(1).$$

Since low types reject $p_1(\lambda)$ and pay at most v_L later, and high types pay at most v_H, $R_1(\lambda) \leq R^c(1)$, which proves (iii).

By Proposition B.2, $\lim \hat{R} < R^c(1)$, so the limit of $\mathcal{D}(\lambda; \hat{R})$ is strictly below $R^c(1)$. Hence, for λ close to one, $R_1(\lambda) \geq \mathcal{D}(\lambda; \hat{R})$ and $\hat{Q} \geq \underline{Q}(\lambda)$. Buyers best respond on path by construction, and off path by Lemma B.4. The seller's first-period offer is optimal, since any other offer yields at most $\mathcal{D}(\lambda; \hat{R})$. Play after rejection for sure by both is the equilibrium $\hat{e}$, and play at every other later history is a continuation equilibrium. Hence the profile is an equilibrium. □