

Multilateral Bargaining with Information Spillovers*

Vladimir Asriyan

William Fuchs

Caio Lorecchio

September 29, 2026

PRELIMINARY DRAFT

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 non-monotone 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 non-monotone 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, Barcelona School of Economics and CEPR (email: vasriyan@crei.cat). Fuchs: University of Texas at Austin and CEPR (email: william.fuchs@mcombs.utexas.edu). Lorecchio: Universitat de Barcelona, BEAT and CREB (email: paeslemelorecchio@ub.edu). 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, 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 and PID2023-148696OA-I00, 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 When valuations are correlated and transactions are observable, a sale to one buyer reveals information about the demand the seller still faces, and therefore about his future pricing incentives. 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. Our key insight is that the information conveyed by trade is *endogenous*. How much information 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 curve can become non-monotone 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 increases with expected early trade by others. By increasing the informational content of each sale, a higher first-period price can, therefore, *increase* rather than decrease the probability of trade. 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,

¹In 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.

information spillovers are strong enough that the feedback between prices and beliefs sustains a continuum of *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): as in Crémer and McLean (1988), learning about one buyer from another reduces information rents, and as correlation approaches one the seller extracts the full surplus, albeit with delay. 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. As correlation rises, the set of equilibrium revenues expands (Proposition 5): its upper bound converges to the commitment benchmark, whereas its lower bound remains bounded away from it. Moreover, this restoration of commitment is not a consequence of the two-period horizon: in an infinite-horizon extension, sufficiently strong information spillovers induce high-valuation buyers to trade almost immediately at prices approaching their valuations, allowing the seller to attain the commitment payoff in the limit. Buyer surplus and total welfare, by contrast, change discontinuously as the equilibrium set changes (Proposition 6). Consequently, revenue, buyer surplus, and welfare can be non-monotone in correlation and need not move together.

These forces have implications for data investments, experimentation, and market design. We show that a seller invests in data—modeled as a costly increase in the correlation of buyers' valuations—only when the investment makes information spillovers sufficiently strong, and that such investment 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 A/B testing, another way by which many modern-day firms acquire information about consumers: by using some buyers as informational probes, the seller refines prices for the remaining buyers, making experimentation profitable only when information spillovers are strong (Corollary 2). A yet different implication of equilibrium multiplicity is that it can mitigate hold-up (Corollary 3). Buyers may coordinate on equilibria that preserve sufficient rents for relationship-specific investments, giving the seller an incentive to facilitate such coordination. This provides a novel rationale for coordination mechanisms, such as supplier associations in Japanese automotive supply chains.

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 et al. \(1986\)](#) 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, non-monotone 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. [Denicolo and Garella \(1999\)](#) show that a durable-goods monopolist may use rationing to mitigate the Coase problem, [Su \(2007\)](#) analyzes dynamic pricing with limited inventory when buyers differ in both their valuations and their costs of waiting, and [Board and Skrzypacz \(2016\)](#) characterize the revenue-maximizing mechanism for a seller facing a random arrival of forward-looking buyers. [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 non-monotone demand and reshapes the seller’s effective commitment power.

Our mechanism is also related to [Asriyan et al. \(2017\)](#), 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 emphasizes dynamic pricing when earlier purchases or consumption experiences convey information about a common payoff-relevant state to subsequent buyers ([Bose et al., 2006](#); [Arieli et al., 2022](#); [Laiho and Salmi, 2026](#)). Another strand emphasizes strategic delays: agents postpone actions to learn from the behavior of others, potentially generating strategic complementarities, equilibrium multiplicity, and inefficient information aggregation ([Chamley and Gale, 1994](#); [Gul and Lundholm, 1995](#); [Chamley, 2004](#); [Rosenberg et al., 2007](#); [Murto and Välimäki, 2013](#); [Achim](#)

and Klein, 2022). Our setting features a different informational channel. Buyers do not learn about an unknown common state from previous actions. Instead, a buyer’s purchase changes the seller’s beliefs about the valuations of other buyers, thus altering continuation prices. The anticipation of this pricing response links buyers’ current purchasing decisions and generates strategic complementarities.

Our analysis also connects to the literature on mechanism design with correlated types. Crémer and McLean (1988) show that a seller can exploit correlation in buyers’ private information to reduce information rents and extract the full surplus. We study a more restrictive environment in which the seller uses posted prices rather than general mechanisms and learns only from buyers’ endogenous purchase decisions. Correlation still allows the seller to condition future prices on information, but because the seller cannot commit, it may either strengthen or weaken his effective commitment power. Thus, correlated information does not simply improve surplus extraction.

Our applications also relate to the literature on dynamic information provision and recommendation (Kremer et al., 2014; Che and Hörner, 2018; Orlov et al., 2020; Bizzotto et al., 2021; Küçükgül et al., 2022; Parakhonyak and Vikander, 2023; Basak and Zhou, 2025), including work on information and price discrimination (Bergemann et al., 2015; Roesler and Szentes, 2017; Bonatti and Cisternas, 2020; Bergemann et al., 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 Tirole, 2000; Fudenberg and Villas-Boas, 2006). Another strand examines the effects of transaction transparency on market outcomes, both theoretically (Madhavan, 1995; Fuchs et al., 2016) and empirically (Goldstein et al., 2007). 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 Appendix.

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:

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

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:

$$\delta^{t-1}p_t. \quad (2)$$

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:

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

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

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

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 $\lambda \rightarrow 1$, 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

²A 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.

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, our main insights do 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 to examine the implications of our theory for policies aimed at consumer data protection.

Strategies and Information Sets. A strategy for the seller maps his information set to a probability distribution over prices. In the first period, his information set is empty. In the second period, his information set 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 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.

³For 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>.

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.⁴

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:

$$\phi(z)\pi_2(z)v_H + (1 - \phi(z))v_L. \quad (5)$$

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 $\pi^* := \frac{v_L}{v_H}$, which is smaller than π_1 by Assumption 1. In particular, from Equation (5):

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

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:

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

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:

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

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,

⁴If a buyer accepted with probability less than one when indifferent, the seller could deviate by slightly lowering the offer.

given buyer symmetry and market anonymity, we can show that all equilibria are symmetric (see Lemma A.1 in Appendix A). 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.⁵ 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 $\pi_1[v_H - \delta(v_H - v_L) - \varepsilon] + (1 - \pi_1)\delta v_L > v_L$. Together with the skimming property, this implies that the L -type never trades in the first period. Hence, it is also without loss to restrict attention to first-period offers 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 $z_g := (g, p_1)$.

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:

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

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:

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

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.

⁵See, for example, Lemma 10.1 in Fudenberg and Tirole (1991).

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 attributable 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 by the logic of 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.⁶ 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

⁶By Lemma 1, which holds regardless of whether the seller can commit, only H -types trade in the first period.

gives commitment its best chance.⁷ We stress that such 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 weakly 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.⁸ 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 news 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 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 $\lambda \rightarrow 1$, 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.

⁷The 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.

⁸By generically unique, we mean that non-uniqueness arises only in knife-edge parameter configurations.

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 non-monotonic (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 non-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 $\tilde{p}_1$, there is an associated trading probability $\sigma \equiv \sigma(\tilde{p}_1)$ 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 $z_b = (b, \tilde{p}_1)$ and $z_g = (g, \tilde{p}_1)$ (see Equation (10)). We now characterize the demand schedule and illustrate its key features in Figure 1.⁹

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 or not, provided they expect a price v_H again in the second period with certainty. This belief is consistent with the equilibrium play only if the seller's posterior belief is sufficiently high, i.e., $\pi_2(z; \sigma) \geq \pi^*$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \leq \underline{\sigma}$ for $\underline{\sigma}$ satisfying:

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

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

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

and where $\rho_H(z_b, \sigma) = 1 - \lambda \sigma$ by Equation (9). Since the seller is indifferent between offering v_H or v_L after *bad news* in the second period, he can mix between them.¹⁰ As the probability $\phi(z_b)$ of the high offer

⁹Unless 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 $\lambda = 0.9$.

¹⁰Since the posterior belief following good news is strictly above π^* , the offer in that event is v_H .

decreases, the seller must decrease the first-period offer to keep the H -types indifferent. The lowest offer $\tilde{p}_1 = \underline{p}$ 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 or not, 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., $\pi_2(z; \sigma) \leq \pi^*$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \geq \bar{\sigma}$ for $\bar{\sigma}$ satisfying:

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

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

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

Since the seller is indifferent between offering v_H or v_L following *good news* in the second period, he can again mix.¹¹ 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 $\tilde{p}_1 = \bar{p}$ arises when v_H is offered with probability one in the second period following *good news*.




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

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

$$\pi_2(z_b; \sigma) < \pi^* < \pi_2(z_g; \sigma), \quad (15)$$

¹¹Since the posterior belief following bad news is strictly below π^* , the offer in that event is v_L .




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

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 $\tilde{p}_1 \in [\underline{p}, \bar{p}]$, the demand schedule is a correspondence that consists of three trading probabilities $\underline{\sigma}$, $\bar{\sigma}$, and $\sigma^*(\tilde{p}_1)$, where $\sigma^*(\tilde{p}_1)$ satisfies the H -type's indifference condition:

$$v_H - \tilde{p}_1 = \delta \rho_H(z_b; \sigma^*(\tilde{p}_1)) (v_H - v_L). \quad (16)$$

We highlight that the non-monotonicity 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 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 non-monotonicity region expands. On the one hand, 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 π^* . 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 π^* .

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 non-monotonic. Specifically, as the first-period price p_1 rises from $\underline{p}$ to $\bar{p}$, the trading probability $\sigma(p_1)$ takes values in $\{\underline{\sigma}, \sigma^*(p_1), \bar{\sigma}\}$, where $\sigma^*(p_1)$ increases from $\underline{\sigma}$ to $\bar{\sigma}$. Moreover, the region of non-monotonicity 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:** $p_1 \in [\underline{p}, \bar{p}]$, $\sigma(p_1) \in \{\sigma^*(p_1), \bar{\sigma}\}$ for $\sigma^*(\cdot)$ defined by Equation (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 $(v_H - \delta(v_H - v_L), \underline{p})$ 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 analogous logic to the one in the previous paragraph, no equilibrium exists with $\sigma < 1$ and $p_1 = v_H - \delta(v_H - v_L)$.

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 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 or 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.¹²

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 $\tilde{p}_1 \neq p_1$. Following such a deviation, both buyers believe that the other buyer will accept the offer with the lowest acceptance probability $\sigma(\tilde{p}_1)$ 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 non-monotonic 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 $v_H - \delta(v_H - v_L)$ —which induces certain first-period trade by high types—or to a price just below $\underline{p}$, which yields trade with probability $\bar{\sigma}$.¹³ As shown in Panel B of the figure, this condition holds

¹²See Sobel and Takahashi (1983), Fudenberg and Tirole (1983), and Fuchs and Skrzypacz (2013).

¹³Following any deviation into 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 receives 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: in the dotted area these equilibria coexist with the high-price equilibrium, while in the striped area they coexist with the low-price equilibrium.

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.

Similarly, the low-price equilibrium exists when its revenue exceeds that from deviating either to 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 the revenue it induces exceeds that from deviating either to $v_H - \delta(v_H - v_L)$, which ensures trade with probability one, or to v_H , which yields trade with probability $\underline{\sigma}$. As shown in the figure, such equilibria arise when information spillovers are sufficiently strong to relax the seller's standard trade-off between charging high prices and sustaining first-period trade. While strategic complementarities among buyers can support equilibria with relatively 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 instead (e.g., point C in Panel A).

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-dominant 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. *When there are multiple equilibria, i.e., $\lambda > \lambda^{nc}$, the Pareto-dominant equilibrium set consists of all contingent-price equilibria with $\sigma = \bar{\sigma}$, and of the low-price equilibrium if the latter exists.*

When an equilibrium is unique, the Pareto-dominant 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 parameter 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 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 $\sigma = \sigma^*(p_1)$ —where $\sigma^*(\cdot)$ 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 better off because he accelerates trade with H -types and trades more frequently with L -types following bad news.

Finally, although contingent-price equilibria with $\sigma = \bar{\sigma}$ dominate those with $\sigma = \sigma^*$, no Pareto ranking exists among contingent-price equilibria with different first-period offers but 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 and low-price equilibria: buyers’ surplus is maximal in the low-price equilibrium, whereas the seller’s revenue is 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:

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

where the news distribution $\{\rho_\theta(z; \sigma)\}$ is given in Equation (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 Equation (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 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 $\bar{R}(\lambda)$ and $\underline{R}(\lambda)$ denote its maximal and minimal elements, respectively.

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

- 1) For $\lambda < \lambda^{nc}$, the unique revenue level $\underline{R}(\lambda) = \bar{R}(\lambda)$ is weakly decreasing in λ .
- 2) For $\lambda \geq \lambda^{nc}$:
 - (i) The lowest attainable revenue, $\underline{R}(\lambda)$, can be non-monotonic in λ and remains bounded away from the seller's payoff under commitment for all λ .
 - (ii) The highest attainable revenue, $\bar{R}(\lambda)$, 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




Figure 4: **Seller's revenue.** The figure depicts the seller's revenue with commitment (x-marked locus) as well as the seller's revenue across all equilibria without commitment (solid line for $\lambda < \lambda^{nc}$ and shaded region between solid loci otherwise). In this figure, we set $\delta = 0.65$.

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 $\lambda \rightarrow 1$, 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 high prices, with $p_1 = v_H$, $\sigma = \underline{\sigma}$, and $\phi(z_g) = \phi(z_b) = 1$ (Figure 4).¹⁴ 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.¹⁵ 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.¹⁶

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.

¹⁴A 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.

¹⁵In this equilibrium, the first-period trading probability $\underline{\sigma}$ decreases with λ (Proposition 2).

¹⁶By contrast, in the other parameter region in which the unique equilibrium features low prices— $p_1 = v_H - \delta(v_H - v_L)$, $\sigma = 1$, and $\phi(z_g) = \phi(z_b) = 0$ —revenue is invariant to λ .

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 $\lambda \rightarrow 1$, revenue in this equilibrium monotonically converges to the commitment payoff.¹⁷

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.¹⁸ Along these equilibria, the effect of λ on seller revenue is non-monotone. 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 ; in our parameterization, it rises).

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 dominant; one revenue-minimizing equilibrium maximizes buyers' surplus, and intermediate revenue levels are also within the Pareto-dominant equilibrium set.

4.2 Spillovers and the Distribution of Surplus

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

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

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

¹⁷The 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 $\bar{\sigma}$ increases with λ (Proposition 2).

¹⁸There 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).




Figure 5: **Social surplus.** The figure depicts the total 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.

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

Proposition 6 (Welfare Set). *The correspondence $\lambda \mapsto \mathcal{W}(\lambda)$ is neither hemi-continuous nor convex-valued. 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, $\bar{W}(\lambda)$, converges to the social surplus under commitment as λ goes to 1.*
- (iii) *The social surplus along any sequence of Pareto-dominant equilibria converges to the social surplus under commitment as λ goes to 1.*

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 equilibrium set changes—either because a new class of equilibria emerges or 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, $\bar{W}(\lambda)$ jumps above $\underline{W}(\lambda)$ at $\lambda = \lambda^{nc}$.

As λ increases further and crosses λ^a , 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. Accordingly, the welfare frontier shifts upward, and $\bar{W}(\lambda)$ jumps at $\lambda = \lambda^a$.

Finally, as λ rises further and crosses λ^b , the low-price equilibrium also disappears, leaving only contingent-price equilibria. At this point, the welfare set contracts discontinuously, with $\bar{W}(\lambda)$ jumping downwards 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 $\lambda \rightarrow 1$. By contrast, along the welfare-minimizing equilibrium—corresponding either to the high-price equilibrium or to the contingent-price equilibrium with intermediate trade σ^* (point B in Panel A of Figure 3)—social surplus remains bounded away from its commitment level. However, if attention is restricted to Pareto-dominant 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, they 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.¹⁹ Analogous definitions apply to buyers and total 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 exists in the absence of information spillovers, that is, when $\lambda = \pi_1$. Then information spillovers benefit the seller if and only if they harm both the buyers and total social surplus.*

¹⁹As multiple equilibria exist when $\lambda > \lambda^{nc}$, it is possible that spillovers benefit the seller in one equilibrium but not in another.

This result follows from the fact that the low-price equilibrium maximizes both buyer surplus and total 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 total 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 exists absent information spillovers, that is, when $\lambda = \pi_1$. Then information spillovers may benefit both the seller and the buyers, or they may reduce both seller revenue and total surplus. However, whenever spillovers benefit the seller, they also benefit the buyers.*

This result is also intuitive. In the high-price equilibrium, buyers receive zero surplus. Thus, 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, buyer surplus and total surplus rise together with seller revenue. By contrast, spillovers can also reduce both seller revenue and total surplus: when λ is small, the seller captures the entire social surplus, and that surplus is locally decreasing in λ , as illustrated in Figure 5.

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, 2019). 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 (European Commission, 2020; Hannak et al., 2014; OECD, 2018; 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 informational spillovers generated by trade. To capture this channel, we allow the seller to invest in data at cost c 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 $\lambda^D > \bar{\lambda}^D$. When investment occurs, buyer surplus and total 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 higher trading probabilities in both periods. By Proposition 8, any data investment that benefits the seller must also increase buyer and total 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 perspective lends support to policies—such as the General Data Protection Regulation (GDPR) or the Digital Services Act (DSA) in the EU—that restrict firms’ ability to collect, store and exploit consumer data.²⁰ 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 et al., 2020). These experiments operate at massive scale: Booking.com reportedly runs more than 1,000 experiments at any time and about 25,000 tests per year, while Expedia conducts concurrent experiments involving millions of users (Thomke, 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.

²⁰See <https://gdpr.eu> and <https://digital-strategy.ec.europa.eu/en/policies/digital-services-act> for details of the GDPR and the DSA, respectively.




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 $\pi_1 = 0.6$ and $\delta = 0.75$, whereas in Panel B, $\pi_1 = 0.5$ and $\delta = 0.7$.

Specifically, the seller offers buyer A a sufficiently low first-period price to induce early trade, while offering buyer B a sufficiently high price that induces delay. Buyer A thus serves as an informational probe, and the seller conditions second-period pricing for B on the information revealed by A 's trading behavior.

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

Experimentation requires the seller to offer buyer A a sufficiently low first-period price, sacrificing revenue on that transaction in exchange for information about buyer B 's valuation. Whether this trade-off is profitable depends on the value of the information generated by buyer A 's purchase decision. 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 . Once correlation exceeds a threshold, however, the seller optimally conditions second-period prices for buyer B on buyer A 's purchase history, making experimentation profitable. 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 in the absence of experimentation. 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 becomes profitable only once the correlation is sufficiently high.

The welfare implications follow the same logic as 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 total 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-valued buyers. Consequently, it destroys gains from trade, so buyer surplus and total 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-valued buyers. Experimentation relaxes these distortions while increasing the seller’s revenue, so both buyer surplus and total 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-valued buyers following good news from buyer *A*, reducing both buyer surplus and total surplus.

While large firms have documented substantial revenue gains from experimentation,²¹ small firms rarely engage in A/B testing despite the decline in experimentation costs. Koning et al. (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 informational spillovers are too weak for experimentation to pay.²²

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 (Grout, 1984; Grossman and Hart, 1986; Hart and Moore, 1990). Apart from vertical integration (Lafontaine and Slade, 2007), 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 (Hines, 1994; 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 informational 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 the hold-up problem is mitigated through repeated interactions that discipline opportunistic

²¹For 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).

²²Although the model features only two buyers, the same force extends naturally to an environment with many buyers (Section 6.2). 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 naturally predicts greater experimentation by larger firms.

behavior. Here, by contrast, no such long-term relationship is required: informational 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 each supplier can make an observable relationship-specific investment at fixed 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 both valuations proportionally, the equilibrium characterization carries over 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, no symmetric equilibrium features investment.*
- (ii) *If $\lambda > \lambda^{nc}$, there exists a threshold $\underline{\alpha}$ such that, whenever $\alpha \geq \underline{\alpha}$, there exists a symmetric equilibrium in which both suppliers invest.*

Consider first the case in which informational 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.

By contrast, when informational 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 this equilibrium, suppliers retain a sufficiently large share of the gains from trade that relationship-specific investment becomes privately profitable whenever α exceeds the threshold $\underline{\alpha}$.

Importantly, the downstream buyer also benefits from coordinating suppliers on this equilibrium. By encouraging *ex ante* investment, the downstream buyer raises suppliers' valuations and thereby the total 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

In this section, we argue that the central mechanism behind our findings extends well beyond the simplifying assumptions adopted in the baseline setting. Specifically, we show that information spillovers continue to generate strategic complementarities in buyers' trading decisions, non-monotone demand, and equilibrium multiplicity in environments with continuous valuations, many buyers, longer trading horizons, and common values. To keep the exposition focused, our discussion here is deliberately high level, with more detailed analysis presented in Appendix B.

6.1 Continuous Valuations

Our baseline model assumes binary valuations, raising the concern that the strategic complementarities and equilibrium multiplicity identified above may be artifacts of the discrete type space and the associated mixed strategies. To address this concern, we extend the analysis to a setting with continuously distributed valuations while preserving the same informational structure. Although this environment is considerably less tractable than our baseline model, we show that the central mechanism continues to operate.

Specifically, we assume that the buyers' valuations are continuously distributed but remain correlated through an underlying aggregate state, so that one buyer's trading decision reveals information about the aggregate state and thus the likely valuation of the other buyer. As in the baseline model, the strength of informational spillovers is governed by a correlation parameter λ . Unlike in the binary model, equilibria are now in pure strategies: for any first-period price, there exists a marginal buyer who is indifferent between purchasing immediately and delaying trade to the second period, with all higher-valuation buyers purchasing in the first period.

The central mechanism behind our main findings remains unchanged. When informational spillovers are weak, a higher first-period price raises the marginal buyer's cutoff valuation and therefore reduces current demand through the standard screening effect. When spillovers are sufficiently strong, however, greater current trade makes favorable information about the remaining buyer more likely, inducing the seller to charge a higher continuation price. Waiting becomes less attractive, lowering the cutoff valuation and increasing current demand. As in the baseline model, buyers' purchase decisions become strategic complements, implying that the demand schedule can be non-monotone and admit multiple pure-strategy equilibria.

Thus, neither strategic complementarities nor equilibrium multiplicity is an artifact of binary valuations or mixed-strategy behavior. The binary model provides a tractable environment in which the mechanism is especially transparent, while the continuous-type extension shows that the same economic force operates in a richer setting. The underlying driver remains the endogenous effect of current trade on the seller's future pricing incentives.

6.2 Many Buyers

Our baseline analysis considers two buyers. We next extend the model to an arbitrary number of buyers and show that the analysis carries over in a straightforward way. The main difference is that the seller's information becomes richer: rather than conditioning on whether a single buyer traded, the seller updates beliefs based on the number of buyers who have purchased. Nevertheless, the underlying mechanism is unchanged. Each buyer's willingness to trade depends on the information generated by the purchases of others, while those purchases are themselves influenced by the seller's continuation pricing incentives. As a result, buyers' purchase decisions remain strategic complements, demand remains non-monotone, and equilibrium multiplicity continues to arise. If anything, these complementarities become stronger as the market grows, because even small changes in other buyers' trading probabilities can substantially affect the information revealed by aggregate trade and hence the seller's future pricing decisions.

6.3 Infinite Horizon

Our baseline model has only two trading opportunities, raising the question of whether information spillovers remain powerful when the seller can repeatedly revise prices. To address this, we extend the model to an infinite horizon. Without information spillovers, the environment reduces to the classic dynamic monopoly problem: buyers delay purchases in anticipation of future price cuts, limiting the seller's ability to sustain high prices.

Whereas frequent opportunities to revise prices exacerbate the seller's commitment problem in the standard Coasian environment, sufficiently strong information spillovers can restore the seller's commitment power. For any fixed discount factor, as buyers' valuations become sufficiently correlated, we construct equilibria in which high-valuation buyers trade almost immediately at prices approaching v_H , while low-valuation buyers trade in the following period at v_L . Consequently, the seller's revenue converges to its commitment counterpart.

The intuition is that an early purchase becomes sufficiently good news about a remaining buyer that the seller subsequently extracts almost all of a high type's surplus. By contrast, following no trade, the seller remains sufficiently pessimistic that he charges v_L in the next period. A high-type buyer who delays therefore faces an increasingly likely continuation in which the seller extracts almost all of her surplus. Anticipating this, high types have little incentive to wait and instead trade immediately at nearly their full valuation.

The same strategic complementarity as in our baseline model underlies this result. More aggressive early trade makes an observed purchase more informative about a remaining buyer, raising the seller's posterior and strengthening his incentive 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. Thus, information spillovers can effectively restore commitment even in an infinite-horizon environment.

6.4 Common Values

We have assumed that a buyer's valuation depends only on her private type or signal θ_i , so trading histories reveal information about the other buyer's valuation but are of no direct payoff relevance for the buyer's own consumption value. We now 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 still privately observes a signal $\theta_i \in \{H, L\}$ at $t = 1$, distributed exactly as in Section 2. The buyer learns her true valuation, however, 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 true valuation is v_{θ_i} only with probability $1 - \gamma$; with the complementary probability it is tied to the other buyer's signal. When $\gamma = 0$, this specification nests our baseline model. For $\gamma > 0$, buyer i 's expected valuation conditional on her own signal is:

$$\mathbb{E}[v_i|\theta_i] = (1 - \gamma)v_{\theta_i} + \gamma \mathbb{E}[v_{\theta_j}|\theta_i], \quad (19)$$

where:

$$\mathbb{E}[v_{\theta_j} | \theta_i = \theta] = \mathbb{P}(\theta_j = H | \theta_i) v_H + (1 - \mathbb{P}(\theta_j = H | \theta_i)) v_L. \quad (20)$$

All equilibrium objects vary continuously in γ , so our qualitative results—presence of strategic complementarities, non-monotone demand, and equilibrium multiplicity—survive also when γ is small. 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. In particular, if one buyer rejects in the first period, the seller infers that this buyer is more likely to have a low signal. Since signals now partially determine both buyers’ realized valuations, the remaining 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, inducing 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 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, although the interaction between these two informational channels is beyond the scope of the present paper.

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 non-monotone, 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.

References

- Achim, P. and Klein, N. (2022). Strategic investment and learning with private information. *Journal of Economic Theory*, 204:105523.
- Adams, J. J., Fang, M., Liu, Z., and Wang, Y. (2025). The rise of AI pricing: Trends, driving forces, and implications for firm performance. *Journal of Monetary Economics*, page 103875.
- Arieli, I., Koren, M., and Smorodinsky, R. (2022). The implications of pricing on social learning. *Theoretical Economics*, 17(4):1761–1802.
- Asriyan, V., Fuchs, W., and Green, B. (2017). Information spillovers in asset markets with correlated values. *American Economic Review*, 107(7):2007–2040.
- Basak, D. and Zhou, Z. (2025). Panics and early warnings. *Journal of Political Economy*, 133(7):2089–2138.
- Bergemann, D., Brooks, B., and Morris, S. (2015). The limits of price discrimination. *American Economic Review*, 105(3):921–957.
- Bergemann, D., Brooks, B., and Morris, S. (2025). On the alignment of consumer surplus and total surplus under competitive price discrimination. *American Economic Journal: Microeconomics*, 17(4):234–259.
- Bizzotto, J., Rüdiger, J., and Vigier, A. (2021). Dynamic persuasion with outside information. *American Economic Journal: Microeconomics*, 13(1):179–194.
- Board, S. and Skrzypacz, A. (2016). Revenue management with forward-looking buyers. *Journal of Political Economy*, 124(4):1046–1087.
- Bonatti, A. and Cisternas, G. (2020). Consumer scores and price discrimination. *The Review of Economic Studies*, 87(2):750–791.
- Bose, S., Orosel, G., Ottaviani, M., and Vesterlund, L. (2006). Dynamic monopoly pricing and herding. *The RAND Journal of Economics*, 37(4):910–928.
- Brynjolfsson, E. and McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5):133–139.
- Brynjolfsson, E. and McElheran, K. (2019). Data in action: data-driven decision making and predictive analytics in us manufacturing. *Working paper*.
- Bulow, J. I. (1982). Durable-goods monopolists. *Journal of Political Economy*, 90(2):314–332.
- Calzolari, G., Felli, L., Koenen, J., Spagnolo, G., and Stahl, K. O. (2021). Relational contracts and trust in a high-tech industry. *CEPR discussion paper No. DP16653*.

- Chamley, C. (2004). Delays and equilibria with large and small information in social learning. *European Economic Review*, 48(3):477–501.
- Chamley, C. and Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica: Journal of the Econometric Society*, pages 1065–1085.
- Che, Y.-K. and Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925.
- Coase, R. H. (1972). Durability and monopoly. *The Journal of Law and Economics*, 15(1):143–149.
- Cr  mer, J. and McLean, R. P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257.
- Denicolo, V. and Garella, P. G. (1999). Rationing in a durable goods monopoly. *The RAND Journal of Economics*, pages 44–55.
- Drugov, M. (2021). Bargaining with informational and payoff externalities. *Journal of Economics & Management Strategy*, 30(2):398–419.
- Dyer, J. H. and Nobeoka, K. (2000). Creating and managing a high-performance knowledge-sharing network: the toyota case. *Strategic Management Journal*, 21(3):345–367.
- Elliott, M., Galeotti, A., Koh, A., and Li, W. (2026). Market segmentation through information. *Review of Economic Studies*, page rdag067.
- European Commission, E. (2020). A european strategy for data. *Strategy report*.
- Fuchs, W.,   ry, A., and Skrzypacz, A. (2016). Transparency and distressed sales under asymmetric information. *Theoretical Economics*, 11(3):1103–1144.
- Fuchs, W. and Skrzypacz, A. (2013). Bargaining with deadlines and private information. *American Economic Journal: Microeconomics*, 5(4):219–243.
- Fudenberg, D. and Tirole, J. (1983). Sequential bargaining with incomplete information. *The Review of Economic Studies*, 50(2):221–247.
- Fudenberg, D. and Tirole, J. (1991). Sequential bargaining under incomplete information. In *Game Theory*, chapter 10, pages 397–434. MIT press.
- Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, pages 634–657.
- Fudenberg, D. and Villas-Boas, J. M. (2006). Behavior-based price discrimination and customer recognition. In Hendershott, T. J., editor, *Handbook on economics and information systems*, chapter 8, pages 377–436. Elsevier.

- Goldstein, M. A., Hotchkiss, E. S., and Sirri, E. R. (2007). Transparency and liquidity: A controlled experiment on corporate bonds. *The review of financial studies*, 20(2):235–273.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4):691–719.
- Grout, P. A. (1984). Investment and wages in the absence of binding contracts: A nash bargaining approach. *Econometrica: Journal of the Econometric Society*, pages 449–460.
- Gul, F. and Lundholm, R. (1995). Endogenous timing and the clustering of agents’ decisions. *Journal of political Economy*, 103(5):1039–1066.
- Gul, F., Sonnenschein, H., and Wilson, R. (1986). Foundations of dynamic monopoly and the coase conjecture. *Journal of economic Theory*, 39(1):155–190.
- Hannak, A., Soeller, G., Lazer, D., Mislove, A., and Wilson, C. (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 conference on internet measurement conference*, pages 305–318.
- Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6):1119–1158.
- Hines, P. (1994). Internationalization and localization of the kyoryoku kai: The spread of best practice supplier development. *The International Journal of Logistics Management*, 5(1):67–72.
- Joskow, P. L. (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. *The American Economic Review*, pages 168–185.
- Kartik, N. and Zhong, W. (2026). Lemonade from lemons: Information design and adverse selection. *Theoretical Economics*, 21(1):281–324.
- Kohavi, R., Tang, D., and Xu, Y. (2020). *Trustworthy online controlled experiments: A practical guide to A/B testing*. Cambridge University Press.
- Kohavi, R. and Thomke, S. (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.
- Koning, R., Hasan, S., and Chatterji, A. (2022). Experimentation and start-up performance: Evidence from A/B testing. *Management Science*, 68(9):6434–6453.
- Kremer, I., Mansour, Y., and Perry, M. (2014). Implementing the “wisdom of the crowd”. *Journal of Political Economy*, 122(5):988–1012.
- Küçükgül, C., Özer, Ö., and Wang, S. (2022). Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Management Science*, 68(7):4899–4918.

- Lafontaine, F. and Slade, M. (2007). Vertical integration and firm boundaries: The evidence. *Journal of Economic Literature*, 45(3):629–685.
- Laiho, T. and Salmi, J. (2026). Coasian dynamics and endogenous learning. *Journal of Economic Theory*, page 106136.
- Levin, J. (2003). Relational incentive contracts. *American Economic Review*, 93(3):835–857.
- Libgober, J. and Mu, X. (2021). Informational robustness in intertemporal pricing. *The Review of Economic Studies*, 88(3):1224–1252.
- Macchiavello, R. and Morjaria, A. (2023). Relational contracts: Recent empirical advancements and open questions. *Journal of Institutional and Theoretical Economics*, 179(3-4):1–28.
- Madhavan, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *The Review of Financial Studies*, 8(3):579–603.
- Murto, P. and Välimäki, J. (2013). Delay and information aggregation in stopping games with private information. *Journal of Economic Theory*, 148(6):2404–2435.
- OECD (2018). Personalised pricing in the digital era. Technical report, Organisation for Economic Co-operation and Development.
- Orlov, D., Skrzypacz, A., and Zryumov, P. (2020). Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578.
- Parakhonyak, A. and Vikander, N. (2023). Information design through scarcity and social learning. *Journal of Economic Theory*, 207:105586.
- Roesler, A.-K. and Szentes, B. (2017). Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–2080.
- Rosenberg, D., Solan, E., and Vieille, N. (2007). Social learning in one-arm bandit problems. *Econometrica*, 75(6):1591–1611.
- Sobel, J. and Takahashi, I. (1983). A multistage model of bargaining. *The Review of Economic Studies*, 50(3):411–426.
- Stokey, N. L. (1979). Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93(3):355–371.
- Su, X. (2007). Intertemporal pricing with strategic customer behavior. *Management Science*, 53(5):726–741.
- Thisse, J.-F. and Vives, X. (1988). On the strategic choice of spatial price policy. *The American Economic Review*, pages 122–137.
- Thomke, S. (2020). Building a culture of experimentation. *Harvard Business Review*, 98(2):40–47.

Appendix A: Proofs for Sections 2.1 -5

Lemma A.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:

$$\pi(z_g; \sigma_i, \sigma_j) = \frac{\lambda[\sigma_i(1 - \sigma_j) + (1 - \sigma_i)\sigma_j]}{\lambda[\sigma_i(1 - \sigma_j) + (1 - \sigma_i)\sigma_j] + (1 - \lambda)(\sigma_i + \sigma_j)}, \quad (\text{A.1})$$

$$\pi(z_b; \sigma_i, \sigma_j) = \frac{\pi_1 \lambda (1 - \sigma_i)(1 - \sigma_j) + \frac{\pi_1(1-\lambda)}{2} [(1 - \sigma_i) + (1 - \sigma_j)]}{\pi_1 \lambda (1 - \sigma_i)(1 - \sigma_j) + \pi_1 (1 - \lambda) [(1 - \sigma_i) + (1 - \sigma_j)] + 1 - 2\pi_1 + \pi_1 \lambda}. \quad (\text{A.2})$$

Let us briefly explain the derivation of these posteriors.

First, note that, in any equilibrium, the low types never trade in the first period (see Lemma 1). Next, since the market is anonymous (Assumption 2), the seller observes only the number of first-period trades. Following good news, exactly one buyer traded. If both buyers have high valuations (which occurs with probability $\lambda\pi_1$), this event occurs with probability $\sigma_i(1 - \sigma_j) + (1 - \sigma_i)\sigma_j$, and the remaining buyer necessarily has a high valuation. If instead exactly one buyer is high-valued (which occurs with probability $\pi_1(1 - \lambda)$), good news occurs only when that high-valued buyer trades, in which case the remaining buyer has a low valuation. Applying Bayes' rule yields $\pi(z_g; \sigma_i, \sigma_j)$ in Equation (A.1).

Following bad news, neither buyer traded and both remain in the market. Conditional on both buyers having high valuations, an arbitrary remaining buyer has a high valuation with probability one. Conditional on exactly one buyer having a high valuation, an arbitrary remaining buyer has a high valuation with probability 1/2. The factor 1/2 reflects anonymity: when exactly one of the two remaining buyers has a high valuation, the seller cannot identify which one. Weighting these events by their probabilities of generating no trade and applying Bayes' rule yields $\pi(z_b; \sigma_i, \sigma_j)$ in Equation (A.2).

We begin the proof 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:

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

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-type buyers are better off by accepting to trade 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-type buyers are better off delaying trades 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 (Equation (A.1)) is strictly higher than that following bad news (Equation (A.2)). The resulting continuation values of high-valuation buyers i and j are, respectively:

$$\begin{aligned} Q_H^i &= \delta [\lambda \sigma_j S_H(z_g) + (1 - \lambda \sigma_j) S_H(z_b)], \\ Q_H^j &= \delta [\lambda \sigma_i S_H(z_g) + (1 - \lambda \sigma_i) S_H(z_b)]. \end{aligned}$$

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 $v_H - p_1 \leq Q_H^j$, so that these inequalities contradict $Q_H^i > Q_H^j$.

Case 2. Suppose that the seller posts a 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:

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

In the former case, the seller can profitably deviate and raise the price slightly without reducing high type's probability of trade (which must equal 1). In the latter, posting a price $p_1 > v_H$ (at which there is no trade in the first period) is weakly dominated by the offer $p_1 = v_H$.

Case 4. Lastly, 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, which requires $\pi_2(z_b; \sigma_i, \sigma_j) \geq \pi^*$. Let $\underline{\sigma}$ denote the symmetric trading probability satisfying:

$$\pi_2(z_b; \underline{\sigma}, \underline{\sigma}) = \pi^*.$$

Since holding the sum $\sigma_i + \sigma_j$ fixed, the posterior $\pi_2(z_b; \sigma_i, \sigma_j)$ as defined in Equation (A.2) is increasing in $(1 - \sigma_i)(1 - \sigma_j)$, we have that $\sigma_i \neq \sigma_j$ implies $(1/2)(\sigma_i + \sigma_j) < \underline{\sigma}$. 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:²³

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

Since the increase in trading probabilities is discrete, the deviation is profitable for ε small enough. $\square$

Proof of Lemma 1. See text. $\square$

Proof of Proposition 1. In what follows, we first solve the seller's problem as if he could choose himself the trading probabilities of both buyers, as well as if he could charge them different prices. Then we show that the unique allocation is the one described in the statement of the proposition. Next, we prove that this unrestricted allocation is implementable as an equilibrium of our restricted commitment benchmark, and it is unique among the seller's preferred equilibria implementing such allocation.

Step 1. Since Lemma 1 applies regardless of whether the seller can commit, the low-valuation buyer never trades in the first period, whereas high-valuation buyer $i \in \{A, B\}$ trade with some probability σ_i . Given any strategy profile for which low-valued buyers never trade in the first period, the seller's revenue from each buyer $i \in \{A, B\}$ is

$$\begin{aligned} R_i = \pi_1 \left\{ \sigma_i p_{1,i} + (1 - \sigma_i) \delta \sum \rho_H(z_i) [\phi_i(z_i) v_H + (1 - \phi_i(z_i)) v_L] \right\} \\ + (1 - \pi_1) \delta \sum \rho_L(z_i) (1 - \phi_i(z_i)) v_L. \end{aligned} \quad (\text{A.3})$$

and buyer i 's maximal acceptable offer in the first period is

$$p_{1,i}^* = v_H - \delta(v_H - v_L) \sum \rho_H(z_i) [1 - \phi_i(z_i)]. \quad (\text{A.4})$$

Since the seller's revenue from each buyer is increasing in $p_{1,i}$, the seller charges each buyer their corresponding maximal acceptable offer. Substituting Equation (A.4) in Equation (A.3) and rearranging yields:

$$R_i = \pi_1 v_H [\delta + (1 - \delta) \sigma_i] + \delta \sum (1 - \phi_i(z_i)) [(1 - \pi_1) v_L \rho_L(z_i) - \pi_1 (v_H - v_L) \rho_H(z_i)]. \quad (\text{A.5})$$

For good news,

$$\rho_H(z_i = z_g) = \lambda \sigma_j, \quad \rho_L(z_i = z_g) = \frac{(1 - \lambda) \pi_1 \sigma_j}{1 - \pi_1}.$$

²³That trading probabilities are symmetric follows by the same argument as in *Case 1*. That they are equal $\underline{\sigma}$ 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 $\phi(z_b) \in (0, 1)$.

Thus, the squared-bracketed term multiplying $1 - \phi_i(z_g)$ is non-positive:

$$(1 - \pi_1)v_L\rho_L(z_g) - \pi_1(v_H - v_L)\rho_H(z_g) = \pi_1\sigma_j(v_L - \lambda v_H) \leq 0.$$

Being so, it is optimal for the seller to set $\phi_i(z_g) = 1$.

For bad news, the squared-bracketed term multiplying $1 - \phi_i(z_b)$ is:

$$w_i(z_b) = v_L - \pi_1 v_H + \pi_1 \sigma_j (\lambda v_H - v_L). \quad (\text{A.6})$$

Thus, it is optimal for the seller to set:

$$\phi_i(z_b) = \begin{cases} 0 & \text{if } w_i(z_b) > 0, \\ 1 & \text{if } w_i(z_b) < 0, \\ \in [0, 1] & \text{if } w_i(z_b) = 0. \end{cases}$$

After optimizing over $\phi_i(z_b)$, $\phi_i(z_g)$, the revenue the seller gets from buyer i (Equation (A.5)) is

$$R_i = \pi_1 v_H [\delta + (1 - \delta)\sigma_i] + \delta \max\{w_i(z_b), 0\}. \quad (\text{A.7})$$

Summing over both buyers, the total revenue is

$$\begin{aligned} R &= \pi_1 v_H [2\delta + (1 - \delta)(\sigma_A + \sigma_B)] + \delta \sum_{i \in \{A, B\}} \max\{w_i(z_b, \sigma_{-i}), 0\}, \\ &= 2\pi_1 \delta v_H + \sum_{i \in \{A, B\}} [\pi_1 v_H (1 - \delta)\sigma_i + \delta \max\{w_i(z_b, \sigma_{-i}), 0\}]. \end{aligned}$$

Revenue is strictly increasing in each trading probability. Indeed, increasing σ_i directly raises the first term associated with buyer i , while

$$\frac{\partial w_j(z_b)}{\partial \sigma_i} = \pi_1 (\lambda v_H - v_L) > 0,$$

so it also weakly raises the continuation revenue obtained from buyer j . Thus, the solution to the relaxed problem has $\sigma_A = \sigma_B = 1$. Substituting $\sigma_j = 1$ into Equation (A.6) gives

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

If $\lambda < \lambda^c$, the relaxed optimum has $\phi_i(z_g) = \phi_i(z_b) = 1$, and Equation (A.4) gives $p_{1,i} = v_H$. If $\lambda > \lambda^c$, it has $\phi_i(z_g) = 1$ and $\phi_i(z_b) = 0$, and $p_{1,i} = v_H - \delta(1 - \lambda)(v_H - v_L)$. Thus, except at $\lambda = \lambda^c$, the relaxed problem has the unique solution described in Proposition 1.

Step 2. It remains to show that the relaxed solution can be implemented under the restrictions of the commitment benchmark. Since the relaxed optimum is symmetric, it requires neither buyer-specific first-period prices nor buyer-specific continuation rules, and hence corresponds to an admissible pricing plan.

Suppose first that $\lambda < \lambda^c$. Consider the plan:

$$p_1 = v_H, \quad \phi(z_g) = \phi(z_b) = 1.$$

An H -type obtains zero surplus from either accepting immediately or waiting, so immediate acceptance by both H -types is an equilibrium of the buyers' game. Low types reject and never trade. Hence, the allocation of the relaxed solution is implementable.

Now suppose that $\lambda > \lambda^c$. Consider instead:

$$p_1 = v_H - \delta(1 - \lambda)(v_H - v_L), \quad \phi(z_g) = 1, \quad \phi(z_b) = 0.$$

If the other H -type accepts in the first period, an H -type who waits obtains expected surplus $\delta(1 - \lambda)(v_H - v_L) = v_H - p_1$. Thus, immediate acceptance by both H -types is again an equilibrium. Low types reject the first-period offer and trade in the second period if and only if bad news occurs. Hence, this plan implements the relaxed allocation.

Under the equilibrium-selection rule of the commitment benchmark, the seller selects the immediate-trade equilibrium whenever the announced plan admits multiple equilibria. Indeed, Step 1 establishes that any allocation with $\sigma_i < 1$ for some buyer yields strictly lower revenue than the relaxed optimum. The two plans above therefore attain the value of the relaxed problem and are optimal in the restricted commitment problem.

Finally, given λ , the corresponding pricing plan in the restricted problem also attains the value of the relaxed problem. Since the relaxed solution is unique whenever $\lambda \neq \lambda^c$, the optimal pricing plan is generically unique. $\square$

Proof of Proposition 2. See text. $\square$

Proof of Proposition 3. See text. $\square$

To prove Theorem 1, we first identify conditions for the existence of each equilibrium class, that is, high-price equilibrium (Lemma A.2), low-price equilibrium (Lemma A.3) and contingent-price equilibria (Lemmas A.4 and A.5). To proceed, it will be useful to define the (per-buyer expected) revenue the seller gets from each equilibrium candidate, following Proposition 3 and Equation (17). The high-price and low-price revenues can be respectively written as:

$$R^{hp} = (1 - \delta)\pi_1 \underline{\sigma} v_H + \delta \pi_1 v_H, \tag{A.8}$$

$$R^{\ell p} = (1 - \delta)\pi_1 v_H + \delta v_L. \tag{A.9}$$

For $p_1 \in [\underline{p}, \bar{p}]$ the contingent-price revenues with $\sigma = \bar{\sigma}$ and $\sigma = \sigma^*(p_1)$ can be respectively written as:

$$R^{cp}(p_1, \bar{\sigma}) = (1 - \delta)\pi_1\bar{\sigma}p_1 + \delta[\pi_1\bar{\sigma}p_1 + (1 - \pi_1\bar{\sigma})v_L], \quad (\text{A.10})$$

$$R^{cp}(p_1, \sigma^*(p_1)) = (1 - \delta)\pi_1\sigma^*(p_1)v_H + \delta[\pi_1\sigma^*(p_1)\lambda v_H + (1 - \pi_1\sigma^*(p_1))v_L]. \quad (\text{A.11})$$

To derive Equation (A.10), we use the fact that the seller is indifferent between charging v_H and v_L following good news when $\sigma = \bar{\sigma}$. To derive Equation (A.11), we use the fact that, when $\sigma = \sigma^*(p_1)$, the seller charges v_H following good news and v_L following bad news. Moreover, from Equation (16),

$$p_1 = v_H - \delta(v_H - v_L)\rho_H(z_b; \sigma^*(p_1)). \quad (\text{A.12})$$

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

Proof. suppose the seller were to deviate from v_H to $\tilde{p}_1$ such that $v_H - \delta(v_H - v_L) \leq \tilde{p}_1 < v_H$. A deviation to $v_H - \delta(v_H - v_L)$ is clearly profitable if $R^{hp} < R^{\ell p}$, and $R^{hp} \geq R^{\ell p}$ rules it out. So consider $\tilde{p}_1 > v_H - \delta(v_H - v_L)$. It cannot be that $\tilde{p}_1 > \underline{p}$. If such deviation was profitable indeed, then there would be some $\phi(z_b)$ such that:

$$v_H - \tilde{p}_1 = \delta(1 - \lambda\underline{\sigma})(1 - \phi(z_b))(v_H - v_L).$$

Because the seller is indifferent following bad news, and the good news belief is always higher than the bad news belief, the seller's revenue would be similar to Equation (A.8), with the difference that v_H is substituted by $\tilde{p}_1$. Clearly, the seller would be always better off by marginally increasing $\tilde{p}_1$ towards v_H .

Consider now $\tilde{p}_1 < \underline{p}$. Then $\bar{\sigma} > \underline{\sigma}$ can be sustained whenever $\phi(z_g)$ is such that:

$$v_H - \tilde{p}_1 = \delta(v_H - v_L)[\lambda\bar{\sigma}(1 - \phi(z_g)) + 1 - \lambda\bar{\sigma}],$$

Since the new trading probability makes the seller indifferent between second period prices following good news, his expected revenue conditional on z_g is still v_L . Following bad news, under the new trading probability, the seller charges v_L for sure. Therefore, the seller's revenue from such deviation is:

$$\pi_1\bar{\sigma}\tilde{p}_1 + (1 - \pi_1\bar{\sigma})\delta v_L.$$

The seller would be better off inducing a higher price. By continuity, if $R^{hp} < R^{cp}(\underline{p}, \bar{\sigma})$, then the seller would find it profitable to deviate to such $\tilde{p}_1$ arbitrarily close to $\underline{p}$. Thus $R^{hp} \geq R^{cp}(\underline{p}, \bar{\sigma})$ prevents such deviation. $\square$

Lemma A.3. *The low-price equilibrium exists if and only if $R^{\ell p} \geq \max\{R^{hp}, R^{cp}(\underline{p}, \bar{\sigma})\}$.*

Proof. That $R^{\ell p} \geq R^{hp}$ is required is clear, so consider $\tilde{p}_1 < v_H$. Suppose the seller was to deviate to $\tilde{p}_1 > \bar{p}$. As before, that would induce H -buyers to trade with probability $\underline{\sigma}$ and there exists $\phi(z_b)$ such that:

$$v_H - \tilde{p}_1 = \delta(1 - \lambda\underline{\sigma})(1 - \phi(z_b))(v_H - v_L).$$

Because the seller is indifferent following bad news, and the good news belief is always higher than the bad news belief, the seller's revenue would be similar to Equation (A.8), with the difference that v_H is substituted by $\tilde{p}_1$. Clearly, the seller would be always better off by marginally increasing $\tilde{p}_1$ towards v_H . $R^{lp} \geq R^{hp}$ rules out such deviation.

Suppose the seller was to deviate to $\tilde{p}_1 \in (\underline{p}, \bar{p})$. Then H -type buyers could respond by trading with probability $\underline{\sigma}$, the lowest consistent trading probability among the contingent-price region. But then raising $\tilde{p}_1$ to v_H would lead to a strictly better payoff than his worst revenue in the contingent-price revenue. Thus, deviations to $\tilde{p}_1 \in (\underline{p}, \bar{p})$ are unprofitable.

Suppose the seller was to deviate to $\tilde{p}_1 \in (v_H - \delta(v_H - v_L), \underline{p})$. As before, H -type buyers would respond by trading with probability $\bar{\sigma}$ and this can be sustained whenever $\phi(z_g)$ is such that:

$$v_H - \tilde{p}_1 = \delta(v_H - v_L)[\lambda\bar{\sigma}(1 - \phi(z_g)) + 1 - \lambda\bar{\sigma}],$$

Since the new trading probability makes the seller indifferent between second period prices following good news, his expected revenue conditional on z_g is v_L . Following bad news, under the new trading probability, the seller charges v_L . Therefore, his revenue from such deviation is $\pi_1\bar{\sigma}\tilde{p}_1 + (1 - \pi_1\bar{\sigma})\delta v_L$. Clearly, the seller would be better off inducing a higher price. By continuity, if $R^{lp} < R^{cp}(\underline{p}, \bar{\sigma})$, then the seller would find it profitable to deviate to such $\tilde{p}_1$ arbitrarily close to $\underline{p}$. $R^{lp} \geq R^{cp}(\underline{p}, \bar{\sigma})$ then prevents such deviation. $\square$

Lemma A.4. For every $p_1 \in [\underline{p}, \bar{p}]$, $R^{cp}(p_1, \sigma^*(p_1)) \leq R^{cp}(p_1, \bar{\sigma})$.

Proof. Consider any $p_1 \in [\underline{p}, \bar{p}]$. The pair $(p_1, \bar{\sigma})$ determines a probability $\phi(z_g) \in (0, 1]$ and $\phi(z_b) = 0$. Moreover, p_1 can be written as:

$$p_1 = v_H - \delta(v_H - v_L)[\rho_H(z_g; \bar{\sigma})(1 - \phi(z_g)) + \rho_H(z_b; \bar{\sigma})]. \quad (\text{A.13})$$

Substituting Equation (A.13) into Equation (A.10) and rearranging, we obtain:

$$R^{cp}(p_1, \bar{\sigma}) = (1 - \delta)\pi_1\bar{\sigma}v_H + \delta \{ [\pi_1\bar{\sigma}[\phi(z_g)\lambda v_H + (1 - \phi(z_g))v_L] + (1 - \pi_1\bar{\sigma})v_L] \}. \quad (\text{A.14})$$

Now, observe that Equations (A.12) and (A.13) are defined for the same p_1 . Thus, it must be that:

$$\phi(z_g) = \frac{\sigma^*(p_1)}{\bar{\sigma}}.$$

Substituting the above equality into Equation (A.14) and rearranging, we obtain:

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

Because $\bar{\sigma} \geq \sigma^*(p_1)$, from the above equation and from Equation (A.11), it is straightforward to check

that $R^{cp}(p_1, \bar{\sigma}) \geq R^{cp}(p_1, \sigma^*(p_1))$. □

Lemma A.5. *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. Note that $R^{cp}(p_1, \bar{\sigma})$ is strictly increasing in $p_1 \in [\underline{p}, \bar{p}]$. Thus, if $R^{cp}(\bar{p}, \bar{\sigma}) < \max\{R^{hp}, R^{\ell p}\}$, no contingent-price equilibrium with $\sigma = \bar{\sigma}$ can exist. From Lemma A.4, no contingent-price equilibrium with $\sigma = \sigma^*(p_1)$ can exist either, as the seller would find it profitable to deviate to either a high price or a low price. From Proposition 3, $\sigma = \bar{\sigma}$ or $\sigma = \sigma^*(p_1)$ are the only possible equilibria contingent-price profiles.

Conversely, suppose $R^{cp}(\bar{p}, \bar{\sigma}) \geq \max\{R^{hp}, R^{\ell p}\}$. The seller has no profitable deviation to $\tilde{p}_1 \in (\bar{p}, v_H]$: that would trigger H -type buyers to trade at $\underline{\sigma}$, for which $\tilde{p}_1 = v_H$ is a best response, which is assumed to be no better than $R^{cp}(\bar{p}, \bar{\sigma})$. Similarly, he has no profitable deviation to $\tilde{p}_1 \in [\underline{p}, \bar{p})$, as the worst continuation value would still happen at $\underline{\sigma}$. Finally, there is no profitable deviation to $[v_H - \delta(v_H - v_L), \underline{p}]$: whenever $\sigma = \bar{\sigma}$, the seller's revenue is increasing in p_1 , so $R^{cp}(\bar{p}, \bar{\sigma})$ yields a strictly higher payoff; at $\tilde{p}_1 = v_H - \delta(v_H - v_L)$ the seller's revenue is weakly lower by assumption.

Finally, that a continuum of contingent-price equilibria exists when $R^{cp}(\bar{p}, \bar{\sigma}) > \max\{R^{hp}, R^{\ell p}\}$ follows from continuity of $R^{cp}(p_1, \bar{\sigma})$ with respect to p_1 . Specifically, it implies the existence of $p' < \bar{p}$ such that $R^{cp}(p_1, \bar{\sigma}) \geq \max\{R^{hp}, R^{\ell p}\}$ for every $p_1 \in [p', \bar{p}]$. Since each such p_1 can be sustained with $\sigma = \bar{\sigma}$ by an appropriate mixing probability following good news, there exists a continuum of seller's mixing up to one. □

Having established Lemmas A.1 to A.5, we proceed to proving Theorem 1.

Proof of Theorem 1. Define:

$$\Delta^{hp} = R^{cp}(\bar{p}, \bar{\sigma}) - R^{hp}, \quad \Delta^{\ell p} = R^{cp}(\bar{p}, \bar{\sigma}) - R^{\ell p}.$$

By Lemma A.5, a contingent-price equilibrium exists if and only if

$$\Delta = \min\{\Delta^{hp}, \Delta^{\ell p}\} \geq 0,$$

and there is a continuum of such equilibria whenever the inequality is strict.

We first show that Δ^{hp} and $\Delta^{\ell p}$ are strictly increasing in λ . Since $\bar{p} = v_H - \delta(1 - \lambda\bar{\sigma})(v_H - v_L)$, straightforward algebra gives:

$$\Delta^{\ell p}(\lambda) = \pi_1 v_H [-(1 - \delta)(1 - \bar{\sigma}(\lambda)) + \delta \lambda \bar{\sigma}^2(\lambda)(1 - \pi^*)], \quad (\text{A.15})$$

$$\Delta^{hp}(\lambda) = \pi_1 v_H (1 - \delta)(\bar{\sigma}(\lambda) - \underline{\sigma}(\lambda)) + \delta v_H [\pi_1 \lambda \bar{\sigma}^2(\lambda)(1 - \pi^*) - (\pi_1 - \pi^*)]. \quad (\text{A.16})$$

Recall from Proposition 2 that $\bar{\sigma}$ is strictly increasing in λ , whereas $\underline{\sigma}$ is strictly decreasing in λ . Hence, both $\Delta^{\ell p}(\lambda)$ and $\Delta^{hp}(\lambda)$ are continuous and strictly increasing in λ . As a consequence, $\Delta(\lambda) = \min\{\Delta^{hp}(\lambda), \Delta^{\ell p}(\lambda)\}$ is continuous and strictly increasing in λ .

On the one hand, when $\lambda = \pi_1$, the two trading thresholds coincide: $\underline{\sigma} = \bar{\sigma} = \sigma^0$, where:

$$\sigma^0 = \frac{\pi_1 - \pi^*}{\pi_1(1 - \pi^*)}.$$

It follows from Equation (A.16) that:

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

On the other hand, as $\lambda \rightarrow 1$, $\bar{\sigma} \rightarrow 1$, and therefore

$$\lim_{\lambda \rightarrow 1} \Delta^{\ell p}(\lambda) = \pi_1 \delta v_H (1 - \pi^*) > 0,$$

while

$$\lim_{\lambda \rightarrow 1} \Delta^{hp}(\lambda) > 0.$$

Thus, $\Delta(\lambda)$ is negative at $\lambda = \pi_1$, and is positive as $\lambda \rightarrow 1$. Hence, there exists a unique $\lambda^{nc} \in (\pi_1, 1)$ such that $\Delta(\lambda^{nc}) = 0$.

Consider $\lambda < \lambda^{nc}$. Then $R^{cp}(\bar{p}, \bar{\sigma}) < \max\{R^{hp}, R^{\ell p}\}$, so Lemma A.5 rules out a contingent-price equilibrium. Moreover, since $R^{cp}(p_1, \bar{\sigma})$ is strictly increasing in p_1 , $R^{cp}(\underline{p}, \bar{\sigma}) < \max\{R^{hp}, R^{\ell p}\}$. If $R^{hp} > R^{\ell p}$, Lemma A.2 therefore implies that the high-price equilibrium exists, while Lemma A.3 rules out the low-price equilibrium. Conversely, if $R^{\ell p} > R^{hp}$, the low-price equilibrium exists and the high-price equilibrium does not. Thus, the equilibrium is generically unique and is either high-price or low-price.

Now consider $\lambda > \lambda^{nc}$. Then $R^{cp}(\bar{p}, \bar{\sigma}) > \max\{R^{hp}, R^{\ell p}\}$, so Lemma A.5 implies that there is a continuum of contingent-price equilibria. It remains to show that these equilibria may coexist with either a high-price or a low-price equilibrium. To see this, note that:

$$R^{hp} - R^{\ell p} = v_H \{-\pi_1(1 - \underline{\sigma}) + \delta [\pi_1(1 - \underline{\sigma}) + \pi_1 - \pi^*]\}.$$

This difference is increasing in $\delta \in (0, 1)$. At $\delta = 0$, it is strictly negative for any λ whereas at $\delta = 1$ it equals $v_H(\pi_1 - \pi^*) > 0$. Since $\underline{\sigma}(\lambda)$ is decreasing in λ , these inequalities continue to hold uniformly in λ for, respectively, sufficiently low and sufficiently high discount. Hence, for sufficiently low δ , at $\lambda = \lambda^{nc}(\delta)$, $R^{cp}(\bar{p}, \bar{\sigma}) = R^{\ell p} > R^{cp}(\underline{p}, \bar{\sigma})$; and for sufficiently high δ , at $\lambda = \lambda^{nc}(\delta)$, $R^{cp}(\bar{p}, \bar{\sigma}) = R^{hp} > R^{cp}(\underline{p}, \bar{\sigma})$.

For sufficiently low δ , because both R^{hp} and $R^{cp}(\underline{p}, \bar{\sigma})$ are continuous in λ , there exists $\varepsilon > 0$ such that, for every $\lambda \in (\lambda^{nc}(\delta), \lambda^{nc}(\delta) + \varepsilon)$:

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

From Lemma A.3 and Lemma A.5, a low-price equilibrium coexists with a continuum of contingent-price equilibria for every such λ .

Analogously, for sufficiently high δ , continuity of revenues with respect to λ implies that:

$$R^{hp} > \max\{R^{\ell p}, R^{cp}(\underline{p}, \bar{\sigma})\}$$

for all $\lambda > \lambda^{nc}(\delta)$ close enough to the threshold. From Lemma A.2 and Lemma A.5, a high-price equilibrium coexists with a continuum of contingent-price equilibria for every such λ . $\square$

Proof of Proposition 4. See Lemma A.4 and text. $\square$

Proof of Proposition 5. Throughout the proof, fix $\delta \in (0, 1)$. We first characterize the lowest and highest attainable revenues. We claim that:

$$\underline{R}(\lambda) = \max \left\{ R^{hp}(\lambda), R^{\ell p}, R^{cp}(\underline{p}, \bar{\sigma}; \lambda) \right\}. \quad (\text{A.17})$$

Indeed, consider any equilibrium yielding revenue R . By deviating to offers arbitrarily close from below to v_H , $v_H - \delta(v_H - v_L)$, and $\underline{p}$, the seller can obtain revenues arbitrarily close to $R^{hp}(\lambda)$, $R^{\ell p}$, and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$, respectively. Hence, the right-hand side of Equation (A.17) is a lower bound for any equilibrium revenue. This bound is attained. If the maximum is $R^{hp}(\lambda)$ or $R^{\ell p}$, this follows directly from Lemmas A.2 and A.3, respectively. Suppose instead that $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) > \max\{R^{hp}(\lambda), R^{\ell p}\}$. Then $(\underline{p}, \bar{\sigma})$ is a contingent-price equilibrium. Indeed, deviations to offers in $[\underline{p}, v_H]$, under the lowest-trade continuation $\underline{\sigma}$ prescribed off path, yield at most $R^{hp}(\lambda)$, whereas deviations below $\underline{p}$ yield strictly less than $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$, with the low-price deviation yielding $R^{\ell p}$. Thus, no deviation is profitable, proving Equation (A.17).

Likewise, we claim that:

$$\bar{R}(\lambda) = \max \left\{ R^{hp}(\lambda), R^{\ell p}, R^{cp}(\bar{p}, \bar{\sigma}; \lambda) \right\}. \quad (\text{A.18})$$

Indeed, the highest revenue among contingent-price candidates is $R^{cp}(\bar{p}, \bar{\sigma}; \lambda)$. Equation (A.10) is increasing in p_1 and, by Lemma A.4, $R^{cp}(p_1, \sigma^*(p_1); \lambda) \leq R^{cp}(p_1, \bar{\sigma}; \lambda)$ for every $p_1 \in [\underline{p}, \bar{p}]$. Attainability then follows from Lemmas A.2, A.3, and A.5.

We next show that $\mathcal{R}(\lambda)$ is convex valued. For $\lambda < \lambda^{nc}$, this follows immediately from equilibrium uniqueness (see Theorem 1). Consider then $\lambda > \lambda^{nc}$. Along the contingent-price branch with $\sigma = \bar{\sigma}$, $R^{cp}(p_1, \bar{\sigma}; \lambda)$ is continuous and strictly increasing in $p_1 \in [\underline{p}, \bar{p}]$, and $(p_1, \bar{\sigma})$ is an equilibrium whenever $R^{cp}(p_1, \bar{\sigma}; \lambda) \geq \max\{R^{hp}(\lambda), R^{\ell p}\}$. Hence, if

$$R^{cp}(\underline{p}, \bar{\sigma}; \lambda) \geq \max\{R^{hp}(\lambda), R^{\ell p}\},$$

the entire branch is sustained in equilibrium. Otherwise, there exists a unique $p^\dagger \in (\underline{p}, \bar{p})$ such that

$$R^{cp}(p^\dagger, \bar{\sigma}; \lambda) = \max\{R^{hp}(\lambda), R^{\ell p}\},$$

and exactly the prices $p_1 \in [p^\dagger, \bar{p}]$ on this branch are sustained. Therefore, the $\bar{\sigma}$ branch alone attains every revenue in

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

Since, by definition, no equilibrium revenue lies outside this interval,

$$\mathcal{R}(\lambda) = [\underline{R}(\lambda), \bar{R}(\lambda)], \quad (\text{A.19})$$

and thus $\mathcal{R}(\lambda)$ is convex valued. The result at $\lambda = \lambda^{nc}$ follows by the same argument, allowing the interval to be degenerate.

We turn to proving that $\lambda \mapsto \mathcal{R}(\lambda)$ is hemi-continuous. The thresholds $\underline{\sigma}(\lambda)$ and $\bar{\sigma}(\lambda)$ are continuous in λ , and hence so are $\underline{p}(\lambda)$, $\bar{p}(\lambda)$, and the revenue functions appearing in Equations (A.17) and (A.18). Since the maximum of finitely many continuous functions is continuous, both $\underline{R}(\lambda)$ and $\bar{R}(\lambda)$ are continuous. Together with Equation (A.19), this implies that $\mathcal{R}(\lambda)$ is hemi-continuous.

We now establish the comparative statics of the two endpoints of the revenue set. Consider first $\lambda < \lambda^{nc}$. In this region the equilibrium is unique and, from theorem 1, $\bar{R}(\lambda) = \underline{R}(\lambda)$ is either $R^{hp}(\lambda)$ or $R^{\ell p}$. Since $R^{hp}(\lambda)$ is decreasing in λ , because $\underline{\sigma}(\lambda)$ is decreasing, whereas $R^{\ell p}$ is independent of λ , $\max\{R^{hp}(\lambda), R^{\ell p}\}$ is weakly decreasing. This proves part (1).

Consider next $\lambda > \lambda^{nc}$. To establish part (2.i), note that $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) = \delta v_L + \pi_1 \bar{\sigma}(\underline{p} - \delta v_L)$, and therefore:

$$\frac{dR^{cp}(\underline{p}, \bar{\sigma}; \lambda)}{d\lambda} = \pi_1 \bar{\sigma}'(\underline{p} - \delta v_L) + \pi_1 \bar{\sigma} \underline{p}'.$$

The first term is strictly positive, since $\bar{\sigma}'(\lambda) > 0$, whereas the second term can have either sign, since $\underline{p}$ can increase or decrease with λ . In particular, as $\delta \rightarrow 0$, $\underline{p} \rightarrow v_H$ and $\underline{p}' \rightarrow 0$, so the derivative above is strictly positive. Thus, $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$ can be increasing in λ , whereas $\max\{R^{hp}(\lambda), R^{\ell p}(\lambda)\}$ is weakly decreasing. Whenever R^{hp} determines $\underline{R}$ over a lower range of correlation and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda)$ determines it at higher levels, $\underline{R}$ first decreases and then increases. The parameterization in Figure 4 provides such an example.

It remains to show that $\underline{R}(\lambda)$ remains bounded away from the seller's payoff under commitment. Let $R^c(\lambda)$ denote the commitment revenue. The proof of Proposition 1 shows that any feasible commitment policy with $\sigma < 1$ is strictly suboptimal. Hence, $R^{hp}(\lambda) < R^c(\lambda)$ and $R^{cp}(\underline{p}, \bar{\sigma}; \lambda) < R^c(\lambda)$ for every $\lambda < 1$, since $\underline{\sigma}, \bar{\sigma} < 1$. Moreover, $R^{\ell p} < \pi_1 v_H \leq R^c(\lambda)$. Thus, Equation (A.17) implies $\underline{R}(\lambda) < R^c(\lambda)$ for every $\lambda < 1$. The gap does not vanish as $\lambda \rightarrow 1$. Let $\underline{\sigma}(1) = \lim_{\lambda \rightarrow 1} \underline{\sigma}(\lambda) < 1$. Since $\bar{\sigma} \rightarrow 1$ and $\underline{p} \rightarrow v_H - \delta(1 - \underline{\sigma}(1))(v_H - v_L)$, we have $R^{\ell p} < R^c(1)$ and

$$\begin{aligned} \lim_{\lambda \rightarrow 1} R^{hp}(\lambda) &= \pi_1 v_H [\delta + (1 - \delta)\underline{\sigma}(1)] < R^c(1), \\ \lim_{\lambda \rightarrow 1} R^{cp}(\underline{p}, \bar{\sigma}; \lambda) &= R^c(1) - \pi_1 \delta(1 - \underline{\sigma}(1))(v_H - v_L) < R^c(1), \end{aligned}$$

where $R^c(1) = \pi_1 v_H + (1 - \pi_1)\delta v_L$. By continuity, $R^c(\lambda) - \underline{R}(\lambda)$ is therefore bounded away from zero.

Finally, we establish part (2.ii). For $\lambda > \lambda^{nc}$, Lemma A.5 implies $\bar{R}(\lambda) = R^{cp}(\bar{p}, \bar{\sigma}; \lambda)$ and $R^{cp}(\bar{p}, \bar{\sigma}; \lambda)$

is equal to $\delta v_L + \pi_1 \bar{\sigma}(\bar{p} - \delta v_L)$. Both $\bar{\sigma}$ and $\bar{p}$ are strictly increasing in λ , and $\bar{p} > \delta v_L$. Hence, $\bar{R}(\lambda)$ is strictly increasing. Moreover, as $\lambda \rightarrow 1$, $\bar{\sigma} \rightarrow 1$ and $\bar{p} \rightarrow v_H$, so:

$$\lim_{\lambda \rightarrow 1} \bar{R}(\lambda) = \pi_1 v_H + (1 - \pi_1) \delta v_L = R^c(1).$$

□

Proof of Proposition 6. Fix $\delta \in (0, 1)$. We first establish that $\mathcal{W}(\lambda)$ is neither convex-valued nor hemi-continuous. Consider a boundary of a coexistence region characterized in Theorem 1, at which a high- or low-price equilibrium and a contingent-price equilibrium yield the same revenue. Their welfare levels W^{cp} and W^j , for $j \in \{hp, lp\}$, are nevertheless different, since Equation (18) implies:

$$W^{cp} - W^j = \pi_1(p_1^j - p_1^{cp}),$$

and $p_1^{cp} \neq p_1^j$. At the boundary at which contingent-price equilibria first appear, no equilibrium attains the welfare levels strictly between these two values. Hence, $\mathcal{W}(\lambda)$ is not convex-valued, and the appearance of the new welfare level implies a failure of lower hemi-continuity. Conversely, at the boundary at which the coexisting high- or low-price equilibrium disappears, its welfare level remains strictly separated from those of the surviving contingent-price equilibria. Its disappearance implies a failure of upper hemi-continuity.

To prove parts (i) – (iii), it is useful to rewrite the social surplus as follows. In any equilibrium, the H -type's indifference condition implies:

$$v_H - p_1 = \delta \sum_{z \in \{z_b, z_g\}} \rho_H(z; \sigma)(1 - \phi(z))(v_H - v_L).$$

Using this condition in Equation (18), together with Equation (17), social surplus can be written as:

$$W(\sigma, \phi; \lambda) = \pi_1[\sigma + (1 - \sigma)\delta]v_H + (1 - \pi_1)\delta \sum_{z \in \{z_b, z_g\}} \rho_L(z; \sigma)(1 - \phi(z))v_L. \quad (\text{A.20})$$

We now prove part (i). For $\lambda > \lambda^{nc}$, consider the interior contingent-price branch of the inverse demand schedule. At $p_1 = \underline{p}$, $\sigma^*(\underline{p}) = \underline{\sigma}$, and

$$R^{cp}(\underline{p}, \underline{\sigma}; \lambda) < R^{hp}(\lambda) \leq \underline{R}(\lambda),$$

whereas at $p_1 = \bar{p}$,

$$R^{cp}(\bar{p}, \bar{\sigma}; \lambda) = \bar{R}(\lambda) > \underline{R}(\lambda).$$

By continuity, and the fact that $R^{cp}(p_1, \sigma^*(p_1); \lambda)$ increases strictly with p_1 , there exists a unique contingent-price equilibrium $(\tilde{p}_1, \sigma^*(\tilde{p}_1))$ satisfying:

$$R^{cp}(\tilde{p}_1, \sigma^*(\tilde{p}_1); \lambda) = \underline{R}(\lambda).$$

From Proposition 5, $\underline{R}(\lambda)$ remains bounded away from the commitment payoff as $\lambda \rightarrow 1$. Hence, $\sigma^*(\tilde{p}_1)$ must remain bounded away from one. Otherwise, along a sequence with $\lambda \rightarrow 1$ and $\sigma^*(\tilde{p}_1) \rightarrow 1$, Equation (A.11) would imply a contradiction:

$$R^{cp}(\tilde{p}_1, \sigma^*(\tilde{p}_1); \lambda) \longrightarrow \pi_1 v_H + (1 - \pi_1) \delta v_L = R^c(1),$$

For an interior contingent-price equilibrium, it must be that $\phi(z_g) = 1$, and $\phi(z_b) = 0$. Thus, Equation (A.20) gives:

$$W^{cp}(\tilde{p}_1, \sigma^*(\tilde{p}_1); \lambda) = \pi_1 [\sigma^*(\tilde{p}_1) + (1 - \sigma^*(\tilde{p}_1)) \delta] v_H + (1 - \pi_1) \delta \left[1 - \frac{\pi_1 (1 - \lambda) \sigma^*(\tilde{p}_1)}{1 - \pi_1} \right] v_L.$$

Since $\sigma^*(\tilde{p}_1)$ remains bounded away from one, $W^{cp}(\tilde{p}_1, \sigma^*(\tilde{p}_1); \lambda)$ remains bounded away from $\pi_1 v_H + (1 - \pi_1) \delta v_L$, which corresponds to the social surplus under commitment as λ goes to 1. Since $\underline{W}(\lambda) \leq W^{cp}(\tilde{p}_1, \sigma^*(\tilde{p}_1); \lambda)$, this proves part (i).

We next prove part (iii). By Proposition 4, every Pareto-dominant equilibrium is either a low-price equilibrium or a contingent-price equilibrium with $\sigma = \bar{\sigma}$. In the former case, Equation (A.20) gives:

$$W^{\ell p} = \pi_1 v_H + (1 - \pi_1) \delta v_L = W^c(1).$$

In the latter case, note that $\bar{\sigma} \rightarrow 1$ as $\lambda \rightarrow 1$ and:

$$\rho_L(z_g; \bar{\sigma}) = \frac{\pi_1 (1 - \lambda)}{1 - \pi_1} \bar{\sigma} \longrightarrow 0.$$

Thus, Equation (A.20) implies that, for any equilibrium mixing $\phi(z_g)$ following good news under $\bar{\sigma}$:

$$W^{cp}(\bar{\sigma}, \phi; \lambda) \longrightarrow \pi_1 v_H + (1 - \pi_1) \delta v_L = W^c(1).$$

Since $W^c(\lambda) \rightarrow W^c(1)$, social surplus along any sequence of Pareto-dominant equilibria converges to the commitment surplus. This proves part (iii).

Finally, part (ii) follows directly from part (iii). □

Proof of Proposition 7. See text. □

Proof of Proposition 8. See text. □

Proof of Corollary 1. See text. □

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 buyer A , the seller's posterior that buyer B is an H -type is $\lambda > \pi^*$. Following no trade, it is:

$$\frac{\pi_1(1 - \lambda)}{1 - \pi_1}.$$

Hence, recalling the proof of Proposition 1, buyer B 's second-period price becomes contingent on buyer A 's trading decision if and only if

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

For $\lambda \leq \lambda^c$, 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

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

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 $\lambda^{A/B} \geq \lambda^c$. Hence, whenever experimentation is profitable, necessarily $\lambda > \lambda^{A/B}$.²⁴

Finally, the intuition for the welfare consequences from A/B testing having either sign are discussed in the text, and the two parameterizations in Figure 6 provide examples in which profitable experimentation respectively increases and decreases buyer surplus and total surplus. $\square$

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, in which suppliers obtain zero surplus. Since investment costs

²⁴The definition does not imply that experimentation is profitable for every $\lambda > \lambda^{A/B}$; the corollary states only a necessary condition. For the parameterizations in Figure 6, the profitable set is an upper interval and $\lambda^{A/B}$ corresponds to the crossing depicted in the figure.

$F > 0$, an investing supplier obtains payoff $-F$, whereas not investing guarantees at least zero. Hence, no symmetric equilibrium features investment.

Now suppose $\lambda > \lambda^{nc}$. Consider a contingent-price equilibrium with $\sigma = \bar{\sigma}$, and let $\underline{\phi}$ denote the minimum probability with which the downstream buyer charges v_H following good news that sustains such an equilibrium. Thus, letting $p_1(\underline{\phi})$ denote the corresponding first-period price, $\underline{\phi}$ is characterized by:

$$R^{cp}(p_1(\underline{\phi}), \bar{\sigma}) = \max \left\{ R^{hp}, R^{\ell p} \right\}.$$

Since investment scales both valuations proportionally, it leaves π^* , $\bar{\sigma}$, and $\underline{\phi}$ unchanged. Hence, if both suppliers invest and coordinate on the contingent-price equilibrium with $\sigma = \bar{\sigma}$ and $\phi(z_g) = \underline{\phi}$, each supplier obtains expected net surplus:

$$-F + \alpha \delta \pi_1 (v_H - v_L) (1 - \lambda \bar{\sigma} \underline{\phi}).$$

Since $1 - \lambda \bar{\sigma} \underline{\phi} > 0$, this payoff increases linearly and without bound in α .

By contrast, following a unilateral deviation to no investment, the deviating supplier's valuation remains bounded by v_H , and therefore so does her expected continuation surplus, independently of α . It follows that there exists a finite $\underline{\alpha} > 0$ such that, for every $\alpha \geq \underline{\alpha}$, investment is optimal for each supplier when the other supplier invests. Thus, a symmetric equilibrium in which both suppliers invest exists. $\square$

Appendix B: Supplementary Results

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

This specification preserves the informational 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 similar strategic complementarities as in the baseline model.

In any symmetric equilibrium, given a first-period price p_1 , there exists a cutoff valuation $v^*(p_1)$ at which buyers are indifferent between trading immediately and delaying trade. 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:

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

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^* 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:

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

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

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

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 7: 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 to accepting a first-period offer p_1 is given by:

$$v^*(p_1) - p_1 = \delta [\rho_{v^*(p_1)}(z_g)S_{v^*(p_1)}(z_g) + (1 - \rho_{v^*(p_1)}(z_g))S_{v^*(p_1)}(z_b)]. \quad (\text{A.21})$$

Equation (A.21) 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 7 plots the inverse demand schedule implied by Equation (A.21). For each first-period price $\tilde{p}_1$, the corresponding value $\sigma(\tilde{p}_1)$ denotes the equilibrium probability that a buyer's valuation exceeds the cutoff $v^*(\tilde{p}_1)$.

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 non-monotone. 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 7, point A and every point on the thick segment B-C of the inverse demand schedule constitute an equilibrium: given buyers' responses to an off-equilibrium price, the seller has no profitable deviation. Hence,

the same mechanism as in the baseline model generates strategic complementarities, non-monotone 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:

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

for every buyer i . Thus, as before, $\lambda = \pi_1$ corresponds to independent buyer types, while correlation increases with λ and becomes perfect as $\lambda \rightarrow 1$.

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 $n \in \{0, \dots, N\}$ 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:

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

where:

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

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

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




Figure 8: The inverse demand schedule and contingent-price equilibrium set (thick segments) with $N = 5$ (Panel A) in comparison to the inverse demand schedule with $N = 50$ (Panel B). In this figure, parameters are same as those used to produce Figure 1.

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 both 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 non-monotone demand schedule, and the same mechanism generates equilibrium multiplicity.

Figure 8 illustrates the inverse demand schedule for a market with $N = 5$ (Panel A) and a market with $N = 50$ (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 8, the equilibrium set lies within the thick segments.²⁵ Thus, the same strategic complementarities as in the baseline model continue to generate non-monotone demand and equilibrium multiplicity in large markets.

B.3 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

²⁵ 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.

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.

In the absence of information spillovers, our environment reduces to the classic dynamic monopoly problem studied by [Fudenberg and Tirole \(1983\)](#) and [Gul et al. \(1986\)](#). Because the seller cannot commit to future prices, buyers have an incentive to delay trade in anticipation of future price reductions, 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 ($\delta \rightarrow 1$), 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 Coasian 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 overturn this classic Coasian logic.

Proposition A.1. *Fix $\delta < 1$. As $\lambda \rightarrow 1$, 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 their commitment counterpart.*

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 $\lambda \rightarrow 1$. 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.

A characterization 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.

Continuation play must therefore keep track of both the number of remaining buyers and the evolution of beliefs along a much richer set of trading histories.

Nevertheless, we conjecture that the strategic complementarities identified above also generate equilibrium multiplicity, as in the baseline model: there exist equilibria in which high-type buyers trade less aggressively, so that early purchases reveal less information and the seller is less able to sustain high continuation prices. The resulting increase in the value of waiting disciplines the seller's initial price, which in turn validates the less aggressive early trade. Thus, although we do not characterize the full equilibrium correspondence, the same feedback between current trade and future pricing that supports the equilibrium in Proposition A.1 should naturally give rise to multiple self-fulfilling trading outcomes.

Proof of Proposition A.1. Fix $\delta < 1$. Let $V(\pi)$ denote the seller's expected discounted revenues 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 and Tirole (1983); Gul et al. (1986)) imply that

$$V(\pi) = v_L \quad \text{and} \quad U(\pi) = v_H - v_L,$$

whenever $\pi < \pi^*$, while

$$V(\pi) \rightarrow v_H \quad \text{and} \quad U(\pi) \rightarrow 0,$$

as $\pi \rightarrow 1$. Consider a first-period strategy in which each high-type buyer accepts the seller's offer with probability $\sigma \in (0, 1)$, while low-type buyers reject. Suppose that σ satisfies

$$\pi_2(z_b; \sigma) < \pi^* < \pi_2(z_g; \sigma),$$

where the posteriors are defined exactly as in our baseline model. Following the bad history z_b , the seller therefore pools in the continuation game by charging v_L . Following the good history z_g , only one buyer remains, and continuation play is given by the equilibrium of the one-buyer game starting from belief $\pi_2(z_g; \sigma)$.

A high-type buyer who rejects the first-period offer obtains the one-buyer continuation payoff $U(\pi_2(z_g; \sigma))$ if the other buyer trades, which occurs with probability $\lambda\sigma$. Otherwise, the continuation history is z_b , and the seller charges v_L in the next period. Hence, the first-period price that makes the high type indifferent between accepting and rejecting satisfies:

$$v_H - p_1 = \delta [\lambda\sigma U(\pi_2(z_g; \sigma)) + (1 - \lambda\sigma)(v_H - v_L)]. \quad (\text{A.22})$$

We now construct a sequence of such offers as $\lambda \rightarrow 1$. There exists a sequence $\sigma(\lambda) \rightarrow 1$ s.t.:

$$\pi_2(z_g; \sigma(\lambda)) \rightarrow 1 \quad \text{and} \quad \pi_2(z_b; \sigma(\lambda)) < \pi^*.$$

Along this sequence:

$$U(\pi_2(z_g; \sigma(\lambda))) \rightarrow 0, \quad 1 - \lambda\sigma(\lambda) \rightarrow 0.$$

It follows from Equation (A.22) that:

$$p_1(\lambda) \rightarrow v_H.$$

Thus, the proposed first-period offer induces high-type buyers to trade with probability converging to one at a price converging to v_H .

Low-type buyers reject in the first period. As $\lambda \rightarrow 1$ and $\sigma(\lambda) \rightarrow 1$, conditional on a buyer being low, the other buyer is low with probability converging to one. The resulting continuation belief following z_b remains below π^* , so the seller charges v_L in the second period and the low type trades. Hence, the revenues generated by the proposed offer converge to:

$$\pi_1 v_H + (1 - \pi_1) \delta v_L,$$

which is the limiting seller payoff under commitment to future prices.

Finally, let the seller choose an optimal first-period offer. The seller can always choose the above-proposed offer, so his equilibrium payoff is bounded below by the revenues it generates. At the same time, his payoff is bounded above by the commitment payoff. Therefore, the seller's equilibrium payoff converges to:

$$\pi_1 v_H + (1 - \pi_1) \delta v_L.$$

Moreover, convergence to this upper bound requires that high-type buyers trade in the first period with probability converging to one and at a price converging to v_H . Any non vanishing probability of delay by a high type, or any price bounded away from v_H , would generate a revenue loss bounded away from zero relative to commitment. Low-type buyers must likewise trade in the following period at v_L with probability converging to one. This establishes the claimed sequence of equilibria. $\square$

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 show that the mechanism generating strategic complementarities, non-monotone demand, and equilibrium multiplicity does not rely on this assumption. In particular, it also operates in the low-prior case, $\pi_1 v_H < v_L$.

In this case, when correlation is sufficiently weak, $\lambda < \hat{\lambda}^{nc}$, the seller optimally trades with both buyer types immediately at price v_L . 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, $\lambda > \hat{\lambda}^{nc}$, the same mechanism as in the baseline model 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 . Thus, as in the baseline model, the inverse demand schedule is non-monotone.

Figure 9 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-value 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 9: 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. We set $\pi_1 = 0.3$ and the rest of the parameters are the same as those used to produce Figures 1 and 2.

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 non-monotone demand, equilibrium multiplicity, and the ability of the seller to sustain higher first-period prices. For the parameterization in Figure 9, 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).