# Multilateral Bargaining with Information Spillovers*

Vladimir Asriyan William Fuchs Caio Lorecchio

July 28, 2026

PRELIMINARY AND INCOMPLETE

**Abstract**

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

*Asriyan: Universitat Pompeu Fabra, Barcelona School of Economics, CREi, and ICREA (email: vasriyan@crei.cat). Fuchs: University of Texas at Austin and CEPR (email: william.fuchs@mccombs.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, and Tsz-Ning Wong for helpful comments and suggestions. We also thank 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, as well as participants at 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. Asriyan acknowledges financial support from the European Research Council (MACFRIC, grant 101229916). 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, 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.

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

Our first main result characterizes equilibria as a function of the strength of information spillovers (Theorem 1). There is a correlation threshold below which the equilibrium is unique and takes one of two familiar forms: a *high-price* equilibrium, in which only high-valuation buyers trade over time, or a *low-price* equilibrium, in which the seller cuts his second-period price to sell to everyone. Above the threshold, information spillovers are strong enough that the feedback between prices and beliefs sustains a continuum of *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 always benefit 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. Buyer surplus and total welfare, by contrast, change discontinuously as the equilibrium set changes (Proposition 6). Consequently, revenue, buyer surplus, and welfare are generally non-monotonic 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 information spillovers are already 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.

Taken together, our analysis shows that the 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 the allocation of surplus.

**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, by contrast, 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 bargaining relationship reveals information about another, leading principals to benchmark agents against one another and giving rise to equilibrium multiplicity. We instead study a monopolist selling to multiple buyers with correlated valuations. Because the monopolist sets prices for 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.

**Organization.** The paper is organized as follows. 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. Section 7 concludes. Formal derivations, proofs and supplementary results are relegated to the Appendix.

## 2 The Model

Consider a monopolist trying to sell a product to two buyers, indexed by $i \in \{A, B\}$, each of whom 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, which we normalize to zero. Buyer $i$ can be one of two types, $\theta_i \in \{H, L\}$, and values the product at $v_{\theta_i}$, 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 $\theta_i$ from purchasing the good at price $p_t$ in period $t \in \{1, 2\}$ is given by:

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

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

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

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

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

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

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

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

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

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

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

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

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

Assumption 1 ensures that asymmetric information is severe enough, so that in a one-shot interaction the seller prefers not to trade with all buyer types. If the seller were to serve both types at the highest price acceptable to both, his payoff would be $v_L$. By contrast, trading only with high types yields an expected payoff $\pi_1 v_H$, which is strictly higher under the assumption. While this assumption simplifies the analysis, 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 be the result of regulation aimed at leveling the playing field among market participants.[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

[2] A correlation level $\lambda$ 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 $\lambda$ approaches one.

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

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.

## 2.1 Preliminaries

It is convenient to highlight some basic properties that all equilibria must satisfy, independent of whether the seller can commit to second-period prices.

**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 as well 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:

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

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

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

If the seller were able to commit to second-period offers, he could choose any $\phi(z) \in [0,1]$. Lack of commitment, however, requires that $\phi(z)$ be optimal for the seller *ex post*. Define $\pi^\star := \frac{v_L}{v_H}$, which is greater than $\pi_1$ by Assumption 1. In particular, from Equation (5):

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

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

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

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

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

where $E_H^i[\cdot]$ indicates that the expectation is taken with respect to the information set of $H$-type buyer $i$. In principle, if buyers were playing asymmetric strategies in the first period, $E_H^i[\cdot]$ could be buyer-specific. However, given buyer symmetry and market anonymity, one 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 also 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, we have that:

**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 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 $v_H - \delta(v_H - v_L) - \varepsilon >$

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

$v_L$, 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 $\varepsilon$ ensure that $\pi_1\left[v_H - \delta(v_H - v_L) - \varepsilon\right] + (1-\pi_1)\delta v_L > v_L$. 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 $\sigma$ of the $H$-type buyer, let $\rho_\theta(z;\sigma) \equiv \mathbb{P}_\sigma(z|\theta)$ denote the $\theta$-buyer's belief that trading history $z \in \{z_b, z_g\}$ is observed before the seller makes his second-period offer. Then:

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

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

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

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

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

### 2.2 Commitment Benchmark

Before turning to the equilibrium analysis, it is useful to consider a benchmark in which the seller can commit to second-period offers. A key feature of this benchmark is that information spillovers always help the seller: they allow him both (i) to realize gains from trade with low-valuation buyers and (ii) to discipline potential delays by high-valuation buyers.

**Proposition 1.** *Suppose the seller can commit to second-period offers. Then an equilibrium exists and is generically unique.*[6] *In it, the H-type buyer always trades in the first period, and there exists a threshold*

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

$\lambda^c \in (\pi_1, 1)$ *such that:*

- ***High price**: if $\lambda < \lambda^c$, the seller offers $v_H$ in both periods and the L-type never trades;*
- ***Contingent price**: if $\lambda > \lambda^c$, the seller offers $v_H$ after good news and $v_L$ after bad news, and the L-type trades if and only if there is bad news.*

Absent commitment, the seller's offers in the second period are constrained. If he becomes too pessimistic, then he must offer a low price $v_L$. Instead, with commitment, second-period offers need not be ex-post optimal, so the seller does not need to adjust continuation offers to trading histories. In particular, for any history $z$, he can choose any probability $\phi(z)$ in Equation (5) even if only low-valuation buyers remain. Hence, as in Bulow (1982) and Stokey (1979), when the buyers' valuations are uncorrelated ($\lambda = \pi_1$), it is optimal for the seller to induce $\sigma = 1$ and set the high price $v_H$ in both periods, resulting in expected revenues of $\pi_1 v_H$. This strategy remains optimal also when correlation is positive but small ($\lambda < \lambda^c$).

When the buyers' valuations are highly correlated ($\lambda > \lambda^c$), the seller can do even better. To understand why, consider the other extreme case, when the correlation is almost perfect ($\lambda \to 1$). Now, the seller sets news-contingent prices: lowers the offer to $v_L$ in the event that no buyer traded in the first period, but keeps it high at $v_H$ if any buyer traded. Because correlation is high, the $H$-type still expects to receive a high offer in the second period; thus, the seller needs to use only a small discount to induce him to trade in the first period. This result echoes Crémer and McLean (1988): buyers' information about one another reduces their rents and strengthens the seller's position. Yet, because the seller is restricted to price posting, information is revealed only through trade, preventing full surplus extraction.

## 3 Equilibrium Characterization

We now turn to our main setting, in which the seller lacks commitment over 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 emerges.

### 3.1 The Demand Schedule

For any first-period offer $\widetilde{p}_1$, there is an associated trading probability $\sigma \equiv \sigma(\widetilde{p}_1)$ with which the $H$-types accept it. This probability defines the seller's second-period posterior belief $\pi_2(z;\sigma)$ following each history $z \in \{z_b, z_g\}$, where $z_b = (b, \widetilde{p}_1)$ and $z_g = (g, \widetilde{p}_1)$ (see Equation (10)). We now characterize the demand schedule and illustrate its key features in Figure 1.[7]

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., $\pi_2(z;\sigma) \geq \pi^\star$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \leq \underline{\sigma}$ for $\underline{\sigma}$ satisfying:

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

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

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

and where $\rho_H(z_b, \sigma) = 1 - \lambda\sigma$ by 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.[8] As the probability $\phi(z_b)$ of the high offer decreases, the seller must decrease the first-period offer to keep the $H$-types indifferent. The lowest offer $\widetilde{p}_1 = \underline{p}$ arises when $v_H$ is offered with probability zero in the second period following *bad news*.

Second, for any first-period offer below $v_H - \delta(v_H - v_L)$, the $H$-type buyers will accept the offer with probability one, since the price will never be below $v_L$ in the second period. When the price is equal to $v_H - \delta(v_H - v_L)$, the $H$-type buyers are indifferent between trading 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 belief is sufficiently low, i.e., $\pi_2(z;\sigma) \leq \pi^\star$ for $z \in \{z_b, z_g\}$, which holds if and only if $\sigma \geq \bar{\sigma}$ for $\bar{\sigma}$ satisfying:

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

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

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

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

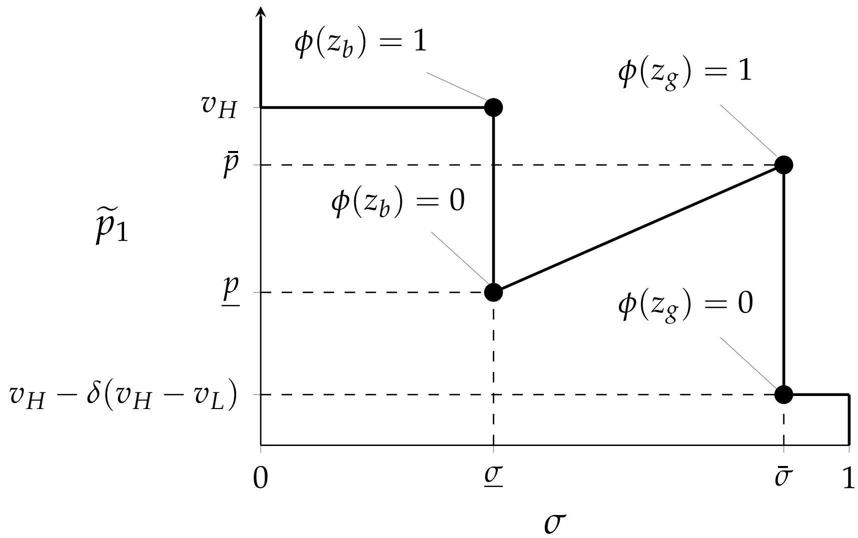


Figure 1: The inverse demand schedule with information spillovers ($\lambda > \pi_1$).

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

Since the seller is indifferent between offering $v_H$ or $v_L$ following *good news* in the second period, he can again mix.[9] As the probability $\phi(z_g)$ of the high offer increases, the seller must increase the first-period offer to keep the $H$-types indifferent. The highest offer $\widetilde{p}_1 = \bar{p}$ arises when $v_H$ is offered with probability one in the second period following *good news*.

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

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

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

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

We highlight that the 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

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

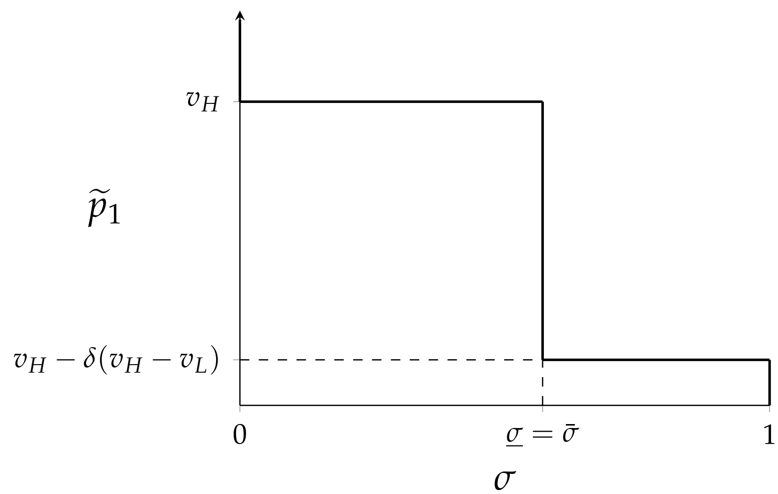


Figure 2: The inverse demand schedule absent information spillovers ($\lambda = \pi_1$).

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 $\lambda$ because higher correlation makes the seller's posterior belief $\pi_2$ *more pessimistic* following *bad news*, requiring high-valuation buyers to trade *less aggressively* in the first period to keep the posterior at $\pi^\star$. On the other hand, $\bar{\sigma}$ increases in $\lambda$ because higher correlation makes the belief $\pi_2$ *more optimistic* following *good news*, requiring high-valuation buyers to trade *more aggressively* in the first period to keep the posterior at $\pi^\star$.

We summarize the key features of the demand schedule discussed above in the following proposition, which we will shortly use 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 $\{\underline{\sigma}, \sigma^\star(p_1), \bar{\sigma}\}$, where $\sigma^\star(p_1)$ 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 $\lambda$, the threshold $\underline{\sigma}$ decreases in it.*

### 3.2 From Demand to Equilibria

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

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

- ***High price:*** *$p_1 = v_H$, $\sigma(p_1) = \underline{\sigma}$, and the L-type never trades.*

- ***Low price:*** *$p_1 = v_H - \delta(v_H - v_L)$, $\sigma(p_1) = 1$, and the L-type trades after every news.*

- ***Contingent price:*** *$p_1 \in [\underline{p},\ \bar{p}]$, $\sigma(p_1) \in \{\sigma^\star(p_1), \bar{\sigma}\}$ for $\sigma^\star(\cdot)$ defined by 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 $\Big(v_H - \delta(v_H - v_L),\ \underline{p}\Big)$ is strictly dominated by a slightly higher offer $\tilde{p}_1$ within the same interval, adjusting $\phi(g, \tilde{p}_1)$ accordingly to keep the first-period trading probability at $\bar{\sigma}$. From Lemma 1, there is no equilibrium in which $p_1 < v_H - \delta(v_H - v_L)$. By 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 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.[10]

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

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

The equilibrium set—illustrated in Figure 3—is constructed as follows. In any equilibrium, the seller posts a first-period price $p_1$, which is accepted with probability $\sigma(p_1)$. For instance, the high-price equilibrium corresponds to $(p_1, \sigma(p_1)) = (v_H, \underline{\sigma})$. Suppose the seller deviates to some offer $\widetilde{p}_1 \neq p_1$. Following such a deviation, both buyers believe that the other buyer will accept the offer with the lowest acceptance probability $\sigma(\widetilde{p}_1)$ consistent with the demand schedule. This off-path belief has bite only for deviations in the interval $[\underline{p}, \bar{p}]$, where inverse demand is 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 $\bar{\sigma}$.[11] As shown in Panel B of the figure, this condition holds when the seller is sufficiently patient: a high discount factor $\delta$ 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

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

[11] Following 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.

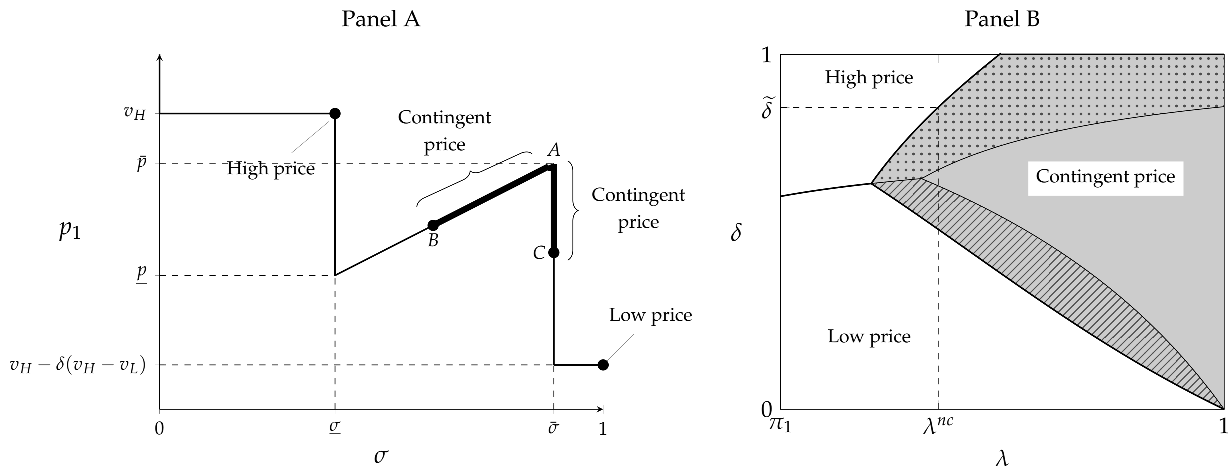


Figure 3: **Equilibrium and coexistence**. Panel A depicts equilibrium prices and trading probabilities for a given set of parameters, whereas Panel B depicts the equilibrium set as a function of $\lambda$ and $\delta$. In Panel B, the gray-shaded area is where contingent-price equilibria exist: in the dotted area these equilibria coexist with the high-price equilibrium, while in the striped area they coexist with the low-price equilibrium.

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 $\delta$ 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 striped-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 $\varepsilon$.

Second, consider the case in which multiple contingent-price equilibria coexist (the shaded region in Panel B). Here, any contingent-price equilibrium with first-period offer $p_1$ and trading probability $\sigma = \sigma^{\star}(p_1)$—where $\sigma^{\star}(\cdot)$ lies on the upward-sloping segment of the demand schedule (see Panel A)—is dominated by another contingent-price equilibrium with the same offer but a higher trading probability $\sigma = \bar{\sigma}$. Buyers earn the same surplus in both equilibria, while the seller is 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 $\sigma = \sigma^{\star}$, 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 her 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 $(p_1, \sigma)$, the seller's per-buyer expected revenue is:

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

where the news distribution $\{\rho_\theta(z;\sigma)\}$ is given in Equation (9). We note that the pair $(p_1, \sigma)$ 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 $\lambda$—and indirectly, through their impact on first-period trade $\sigma$ and price $p_1$, and the seller's ability to condition continuation prices on observed news.

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

**Proposition 5** (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 $\underline{R}(\lambda) = \overline{R}(\lambda)$ is weakly decreasing in $\lambda$.*

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

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

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

Figure 4 illustrates Proposition 5 and contrasts it with the full-commitment benchmark. Under full commitment to second-period offers—our benchmark in Section 2.2—spillovers unambiguously benefit the seller, as shown by the x-marked locus in the figure. As Proposition 1 establishes, when

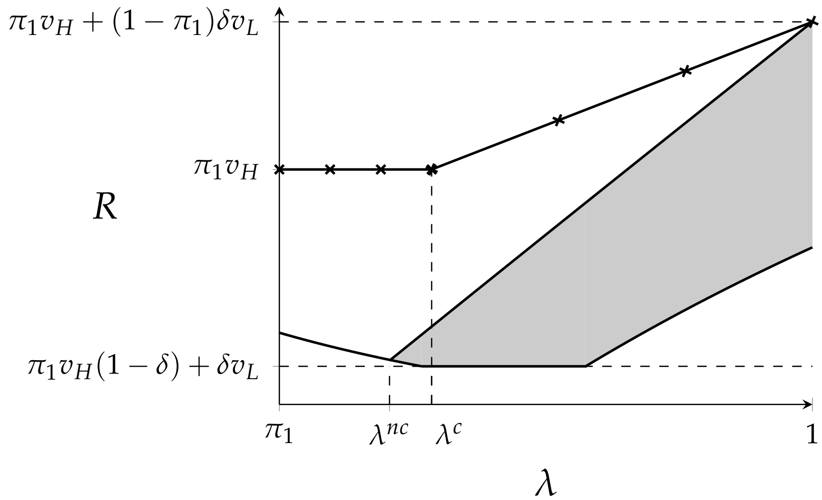


Figure 4: **Seller's revenue**. The figure depicts the seller's revenue with commitment (x-marked locus) as well as the seller's revenue across all equilibria without commitment (solid line for $\lambda < \lambda^{nc}$ and shaded region between solid loci otherwise). In this figure, we set $\delta =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

$\lambda < \lambda^c$ the committed seller's payoff is invariant to spillovers, since he optimally disregards the news when setting prices. Once $\lambda$ exceeds $\lambda^c$, the seller conditions his second-period offer on the news realization, thereby imposing less favorable continuation terms on the high type. As correlation increases, this continuation threat becomes more effective at disciplining delay, enabling the seller to extract a larger share of surplus in the first period. In the limit as $\lambda \to 1$, the seller extracts the full surplus, albeit with delay.

Absent commitment, the effects of information spillovers change markedly, and it is useful to distinguish between two cases, depending on whether the equilibrium is unique. Suppose first that spillovers are weak, so that $\lambda < \lambda^{nc}$ and the equilibrium is unique. Consider the parameter region in which equilibrium play features high prices, with $p_1 = v_H$, $\sigma = \underline{\sigma}$, and $\phi(z_g) = \phi(z_b) = 1$ (Figure 4).[12] In this region, revenue is decreasing in $\lambda$. 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.[13] 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.[14]

If correlation increases beyond $\lambda^{nc}$, however, a new class of equilibria with contingent pricing emerges (Theorem 1). In these equilibria, the seller conditions second-period prices on realized

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

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

[14]By 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 $\lambda$.

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* commitment.

Now suppose $\lambda > \lambda^{nc}$, so that there are multiple equilibria (Theorem 1). In this region, comparative statics are equilibrium-dependent. Along the revenue-maximizing equilibrium, the seller sets contingent second-period prices. Because he already charges a low price following bad news, a further increase in $\lambda$ does not erode his ability to sustain a high price after good news. Moreover, as $\lambda$ increases, a high-type buyer assigns a higher probability to the other buyer also being a high type and therefore anticipates a greater likelihood of facing a high second-period price. This relaxes the first-period incentive constraint: high types are willing to accept a higher initial price, allowing the seller to raise $p_1$ and increase revenue (as shown by the upper boundary of the shaded region in Figure 4). As $\lambda \to 1$, revenue in this equilibrium monotonically converges to the full-commitment payoff.[15]

By contrast, along revenue-minimizing equilibria, spillovers may further *undermine* 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.[16] Along these equilibria, the effect of $\lambda$ on seller revenue is non-monotone. Revenue decreases with $\lambda$ when the high-price equilibrium exists (low $\lambda$ in Figure 4), is constant in $\lambda$ when the low-price equilibrium exists (intermediate $\lambda$), and may either rise or fall with $\lambda$ when neither the high- nor the low-price equilibrium is sustainable (high $\lambda$; 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 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.

[15] The revenue-maximizing equilibrium features $p_1 = v_H - \delta\rho_H(z_b;\bar{\sigma})(v_H - v_L)$, $\sigma = \bar{\sigma}$, $\phi(z_g) = 1$, and $\phi(z_b) = 0$, where $\bar{\sigma}$ increases with $\lambda$ (Proposition 2).

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

## 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:

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

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.

Given correlation $\lambda$, 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 full commitment as $\lambda$ goes to* 1*.*

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

*(iii) The social surplus along any sequence of Pareto-dominant equilibria converges to the social surplus under full commitment as $\lambda$ 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 $\lambda$, whereas the revenue set evolves continuously. These discontinuities arise at values of $\lambda$ where equilibrium selection 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 her 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 $\lambda$ crosses the threshold $\lambda^{nc}$, 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

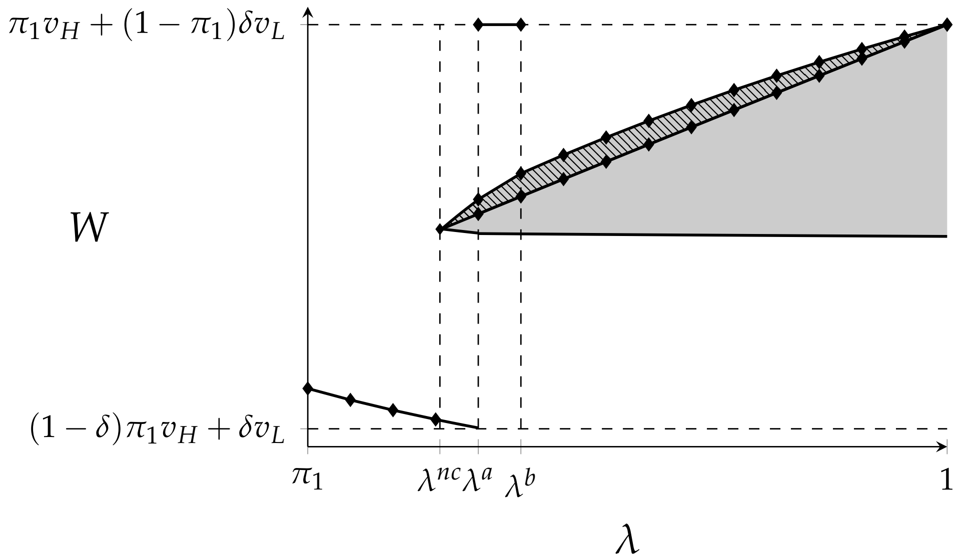


Figure 5: **Social surplus**. The figure depicts the total social surplus across all equilibria without commitment (solid line for $\lambda < \lambda^{nc}$ and shaded region between solid loci otherwise). The figure uses the same parameters as Figure 4.

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 $\lambda$ increases further and crosses $\lambda^a$, 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 $\lambda$ rises further and crosses $\lambda^b$, 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 $\lambda$, only contingent-price equilibria exist and the welfare set evolves continuously with $\lambda$.

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 full-commitment benchmark as $\lambda \to 1$. By contrast, along the welfare-minimizing equilibrium—corresponding either to the high-price equilibrium or to the contingent-price equilibrium with intermediate trade $\sigma^\star$ (point $B$ in Panel A of Figure 3)—social surplus remains bounded away from its full-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 full-commitment counterparts as $\lambda$ 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 equilibrium selection.

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 $\lambda$ benefit (harm) the seller if revenue $R(p_1, \sigma; \lambda)$ is higher (lower) than in the benchmark.[17] 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 information 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.*

This result follows from the fact that the low-price equilibrium maximizes both buyer surplus and total surplus among all equilibria for any $\lambda$. 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, 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 $\lambda$ is small, the seller captures the entire social surplus, and that surplus is locally decreasing in $\lambda$, as illustrated in Figure 5.

As we discuss next, these results provide insights into why firms may have incentives 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,

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

as it may strengthen information spillovers and shift equilibrium selection 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 $\pi_1$ 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.[18] 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. 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 buyer $B$ 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

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

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 $\lambda$, 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,[19] small firms rarely engage in A/B testing despite the dramatic 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 often 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.[20]

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

[20] Although the model features only two buyers, the same force extends naturally to larger markets (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.

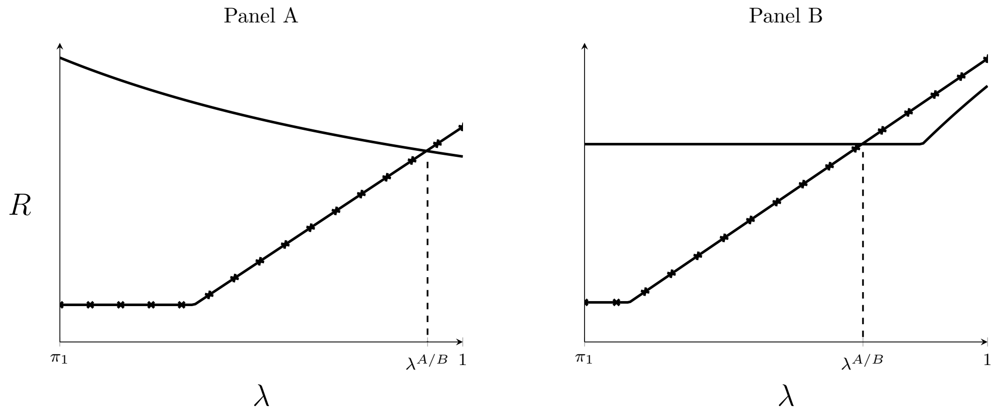


Figure 6: **Revenue with and without experimentation**. The figure depicts the seller's total revenue without A/B testing (solid lines) and with A/B testing (x-marked lines). In both panels, $v_L = 0.8$ and $v_H = 2$. In addition, in Panel A, we set $\pi_1 = 0.6$ and $\delta = 0.75$, whereas in Panel B, $\pi_1 = 0.5$ and $\delta = 0.7$.

## 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 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 $\alpha v_\theta$. 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 $\delta$ is sufficiently high, no symmetric equilibrium features investment.*

*(ii) If $\lambda > \lambda^{nc}$, there exists a threshold $\underline{\alpha} > 0$ 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 $\delta$ 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 $\alpha$ 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 show 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 $\lambda$. 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 therefore 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 assumes two trading periods. One may therefore wonder whether our findings survive when buyers have many opportunities to delay trade. To address this concern, we extend the model to an infinite-horizon environment.

In the absence of informational spillovers, the extension reproduces the standard Coasian logic. Equilibrium features a declining sequence of prices over time, with buyers delaying purchases in anticipation of future price cuts. As the discount factor $\delta$ converges to one—equivalently, as the time between trading opportunities vanishes—the initial price converges to $v_L$, eliminating the seller's rents in the limit. Thus, absent informational spillovers, the model exhibits the familiar Coase conjecture.

We show that informational spillovers fundamentally alter this conclusion. For any fixed discount factor $\delta < 1$, as the correlation parameter $\lambda$ converges to one, there exists a sequence of equilibria in which the high-value buyer purchases in the first period with probability arbitrarily close to one at a price arbitrarily close to $v_H$, while the low-value buyer purchases in the second period with probability arbitrarily close to one at price $v_L$. The intuition mirrors that of the baseline model. Along this sequence, a first-period sale (good news) leads the seller to infer with near certainty that the remaining buyer has a high valuation, sustaining a high continuation price, whereas the absence of trade (bad news) induces the seller to offer $v_L$ immediately. Anticipating that waiting is unlikely to lead to a lower future price, high-value buyers optimally purchase early despite the seller's inability to commit.

Consequently, the seller's expected revenue converges to the full-commitment benchmark. Informational spillovers therefore continue to discipline buyers' incentives to delay by making early trade sufficiently informative about future pricing. Although we do not characterize the full equilibrium set in this richer setting, we show that the same strategic complementarities that drive our baseline results remain operative in the infinite-horizon environment and continue to generate non-monotone demand and equilibrium multiplicity.

### 6.4 Common Values

We have assumed that a buyer's valuation depends only on her private type or signal $\theta_i$, 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 $v_{\theta_i}$, and with probability $\gamma \in [0, \frac{1}{2})$ it instead equals $v_{\theta_j}$. Thus, $\theta_i$ remains the buyer's private information at the trading stage, but her true valuation is $v_{\theta_i}$ 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:

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

where:

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

All of our equilibrium objects vary continuously in $\gamma$, so our qualitative results—presense of strategic complementarities, non-monotone demand, and equilibrium multiplicity—survive also when$\gamma$ is small. When $\gamma$ 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 interdependent 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, 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.

Crémer, 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., Öry, 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 Rconomy*, 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 âwisdom of the crowdâ. *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 3-5 (INCOMPLETE)

**Lemma A.1.** *In equilibrium, type-H buyer i and j accept the first-period offer with the same probability $\sigma$.*

*Proof.* We only need to show that the first-period trading probabilities of high-valuation buyers $i$ and $j \neq i$ are equal, $\sigma_i = \sigma_j$. Suppose to the contrary that $\sigma_i > \sigma_j$.

First, suppose that the seller posts an offer:

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

Second-period prices must then be contingent on news, so that:

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

The continuation values of high-valuation buyers $i$ and $j$ are therefore:

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

and

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

respectively. Since $\sigma_i > \sigma_j$ and $S_H(z_g) < S_H(z_b)$, it follows that:

$$Q_H^i > Q_H^j.$$

Moreover, $\sigma_i > 0$ requires:

$$v_H - p_1 \geq Q_H^i,$$

whereas $\sigma_j < 1$ requires:

$$v_H - p_1 \leq Q_H^j.$$

These inequalities contradict $Q_H^i > Q_H^j$.

Second, suppose that the seller posts:

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

Since $\sigma_j < 1$, the seller can profitably 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.

Third, it is straightforward to rule out equilibrium offers satisfying:

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

In the former case, the seller can profitably deviate and raise the price slightly without reducing 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 $p_1 = v_H$.

Finally, suppose that the seller posts $p_1 = v_H$. 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:

$$\pi_2(z_b) \geq \pi^\star.$$

Under asymmetric trading probabilities, the seller's posterior following bad news is:

$$\pi_2(z_b) = \frac{\lambda\pi_1(1-\sigma_i)(1-\sigma_j)}{\lambda\pi_1(1-\sigma_i)(1-\sigma_j) + \left(1 - \frac{1-\lambda}{1-\pi_1}\pi_1\right)(1-\pi_1)}.$$

Let $\underline{\sigma}$ denote the symmetric trading probability satisfying:

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

Since $\pi_2(z_b)$ is increasing in $(1-\sigma_i)(1-\sigma_j)$, the requirement $\pi_2(z_b) \geq \pi^\star$, together with $\sigma_i \neq \sigma_j$, implies that:

$$\frac{\sigma_i + \sigma_j}{2} < \underline{\sigma}.$$

The seller can therefore profitably lower the first-period offer slightly, inducing the buyers to trade with the symmetric trading probability $\underline{\sigma}$ in the continuation game. The resulting discrete increase in the average first-period trading probability, from $(\sigma_i + \sigma_j)/2$ to $\underline{\sigma}$, dominates the arbitrarily small reduction in the first-period offer, contradicting seller optimality.

We conclude that $\sigma_i = \sigma_j$ in every equilibrium. □

*Proof of Lemma 1.* See text. □

*Proof of Proposition 1.* Since Lemmas 1 and A.1 apply regardless of whether the seller can commit, the low-valuation buyer never trades in the first period, whereas the high-valuation buyer trades with some probability $\sigma$.

We first show that $\sigma = 1$. Suppose instead that $\sigma < 1$. Since the seller can commit to arbitrary second-period offers, he can always increase the continuation payoff of the high-valuation buyer while simultaneously increasing the first-period price so as to keep that buyer indifferent. This allows the seller to induce $\sigma = 1$ without reducing his revenue from the high-valuation buyer. Because trade with the high type is accelerated from the second to the first period with strictly positive probability, the seller additionally gains from discounting. Hence, any equilibrium with $\sigma < 1$ is strictly dominated by one with $\sigma = 1$.

Given $\sigma = 1$, there are only two possible continuation policies:

1. charge $v_H$ after every history; or
2. charge $v_H$ after good news and $v_L$ after bad news.

Under the first policy, expected revenue equals

$$\pi_1 v_H.$$

Under the second policy, the first-period price satisfies:

$$v_H - p_1(\lambda) = \delta(1-\lambda)(v_H - v_L),$$

so expected revenue is:

$$\pi_1 p_1(\lambda) + \delta(1-\pi_1)v_L.$$

Therefore, the contingent-pricing policy is optimal if and only if:

$$\pi_1 v_H \leq \pi_1 p_1(\lambda) + \delta(1-\pi_1)v_L.$$

This inequality defines the threshold $\lambda^c \in (\pi_1, 1)$. ☐

*Proof of Proposition 2.* See text. ☐

*Proof of Proposition 3.* See text. ☐

*Proof of Theorem 1.* ☐

*Proof of Proposition 4.* ☐

*Proof of Proposition 5.* ☐

*Proof of Proposition 6.* ☐

*Proof of Proposition 7.* ☐

*Proof of Proposition 8.* ☐

*Proof of Corollary 1.* ☐

*Proof of Corollary 2.* $\square$

*Proof of Corollary 3.* $\square$

# Appendix B: Supplementary Results (INCOMPLETE)

## 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 $\pi_1$, 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 $\lambda \to 1$, 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 $\lambda$. The key mechanism is therefore unchanged. A buyer's incentive to delay trade depends on the information that other buyers' trading decisions reveals 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 $v^\star(p_1)$ 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:

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

is the probability of news $z$ conditional on aggregate state $\omega$. 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 to trade in that period.

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

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

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

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

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

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

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

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

Equation (A.1) 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.1). For each first-period price $\widetilde{p}_1$, the corresponding value $\sigma(\widetilde{p}_1)$ denotes the equilibrium probability that a buyer's valuation exceeds the cutoff $v^\star(\widetilde{p}_1)$.

As in the baseline model, sufficiently strong correlation generates an upward-sloping region of the inverse demand schedule. The intuition is identical. When buyers expect more first-period trade, observing another buyer purchase becomes more likely. This makes the seller more optimistic about the valuation of buyers who delay, inducing a higher continuation price. Since waiting becomes less attractive, buyers become more willing to purchase immediately. Consequently,

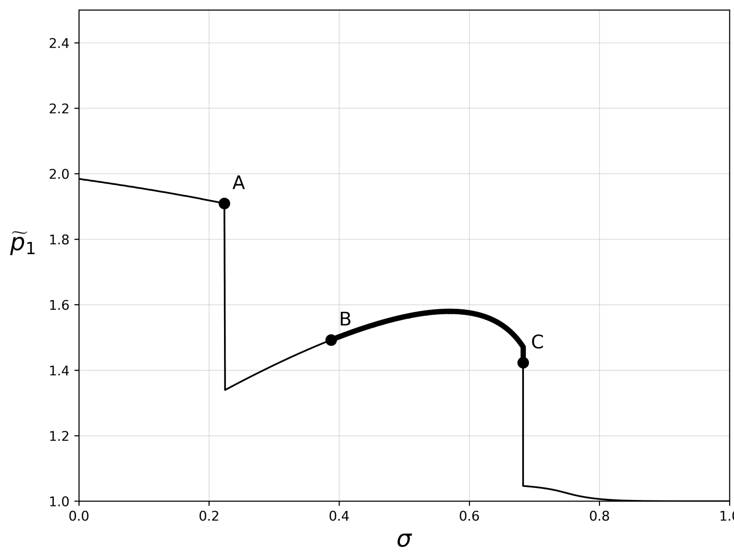


Figure 7: The inverse demand schedule with continuous valuations. In this figure, the parameters used are: $\pi_1 = 0.7$, $v_L = 1$, $v_H = 2.5$, $\delta = 0.95$, and $\lambda = 0.95$.

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:

$$\mathbb{P}(\theta_i = H|\omega = h) = \lambda$$

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

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 $n \in \{0, \ldots, N\}$ of these buyers trade. If in equilibrium a high-valuation buyer accepts the first-period offer with probability $\sigma$, then conditional on buyer $i$'s type, the probability of observing the news realization $z_n$ is:

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

where:

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

is the probability that any one of the remaining buyers trades in the first period conditional on the aggregate state. As in the baseline model, the news distribution is endogenous because it depends on the trading probability $\sigma$, 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:

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

The key difference relative to the baseline is that the seller's first-period price now affects buyers' continuation values through a richer distribution of future news realizations. By changing the first-period trading probability $\sigma$, 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

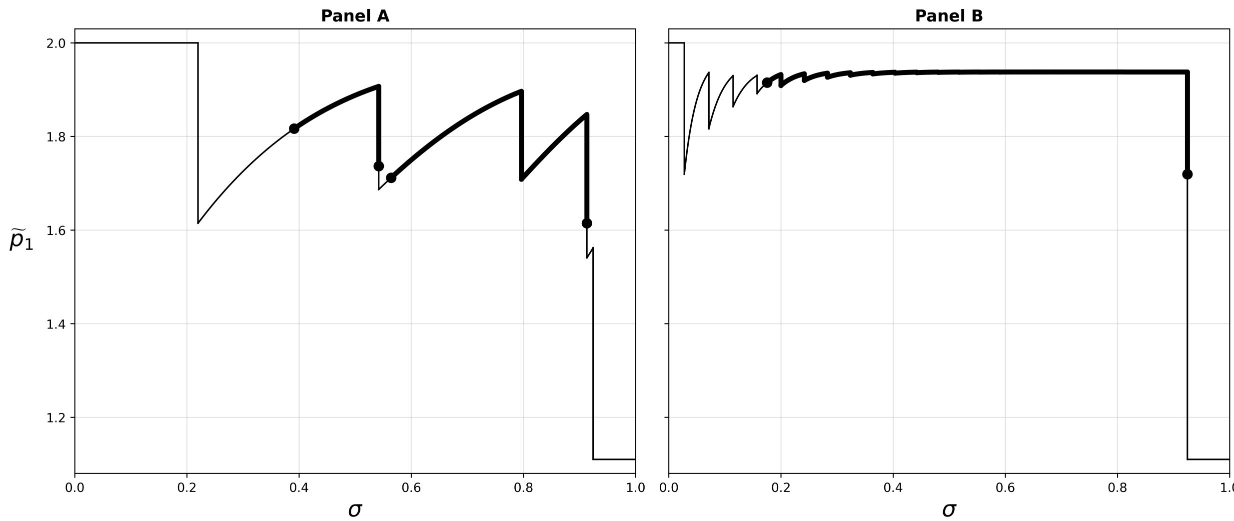


Figure 8: The inverse demand schedule and contingent-price equilibrium set (thick segments) with $N = 5$ (Panel A) in comparison to the inverse demand schedule with $N = 0$ (Panel B). In this figure, parameters are same as those used to produce Figure 1.

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, the equilibria can be computed numerically. For the parameterization in Figure 8, every point on the thick shaded segment of the inverse demand schedules constitutes an equilibrium.[21] 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

## 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, $\lambda < \widehat{\lambda}^{nc}$, 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, $\lambda > \widehat{\lambda}^{nc}$, the same

[21] 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.

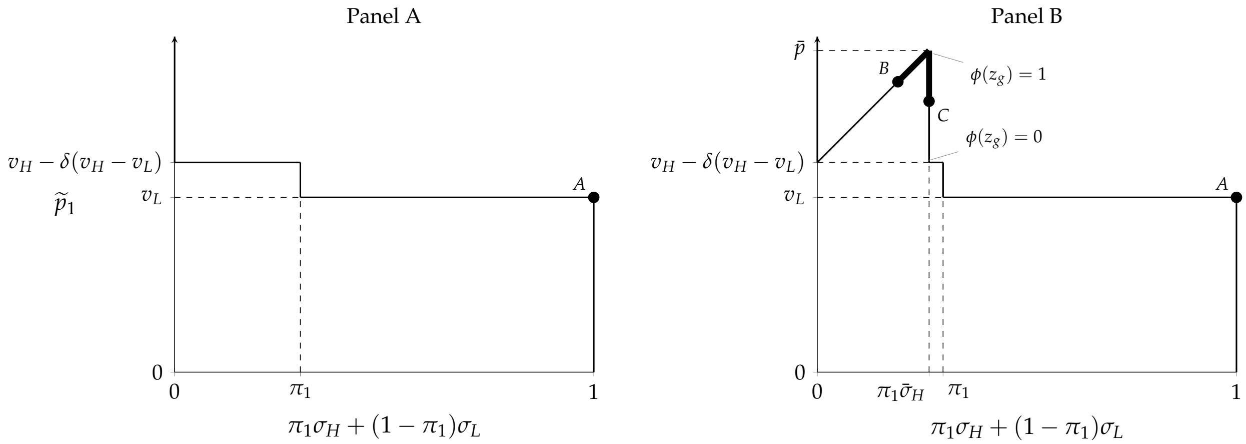


Figure 9: The inverse demand schedule in the low-prior case. Panel A depicts the case of no information spillovers, whereas Panel B depicts the case with information spillovers. We set $\pi_1 = 0.3$ and the rest of the parameters are the same as those used to produce Figures 1 and 2.

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 weakly downward sloping.

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 $\sigma_H$ 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 $p_1 < v_L$.

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 $A$).