

Trade Delay in Bargaining over Lemons under Second-Order Uncertainty

October 22, 2026

Abstract

A seller holds a good that may be high or low quality and proposes prices over time to a buyer who receives an exogenous imperfect private message about the quality before bargaining begins. In a benchmark with an uninformed seller, the Coase conjecture emerges and is driven by two key properties: the market-clearing price is invariant to the seller's deviation, and the seller becomes increasingly pessimistic after repeated rejections. By contrast, when the seller privately observes quality, prices serve not only as offers but also as signals that affect the buyer's expected valuation. When no message can reveal the seller to be of the low type, deviations can induce sufficiently pessimistic off-path beliefs, so the off-path market-clearing price falls below the on-path one. This breaks the invariance of the market-clearing price and makes early market clearing suboptimal. When some message reveals the seller to be of the low type, the high-type seller knows with certainty that the buyer is optimistic; this breaks the seller-pessimism property and allows the high type to separate through delay. Both cases lead to a breakdown of the Coase conjecture and to trade delay in the limit, but differ in the strength of persistence: in the former, delay is merely sustainable, whereas in the latter it is uniformly bounded away from zero. The paper thus identifies two distinct channels through which the interaction between lemons-market valuation and second-order uncertainty breaks the Coase conjecture.

Keywords— Bargaining; Trade Delay; Second-Order Uncertainty; Coase Conjecture; Lemons Market

1 Introduction

Consider a bargaining process in a second-hand car market: a seller (he) attempts to sell a good to a buyer (she). The seller is privately informed about the quality of the car. Before bargaining begins, the buyer receives imperfect information about the quality of the car from previous customers. However, since such information is unverifiable to the seller, the buyer cannot use it as evidence to negotiate a lower price. Instead, she must keep this information private during the bargaining process. On the one hand, although the seller directly cares only

about the payment he receives, he understands that his reputation depends on the true quality of the car and matters to the buyer. Therefore, he uses his private information about the car's quality to infer the buyer's belief and choose prices accordingly. On the other hand, the buyer may be able to leverage her private information to obtain information rents.

In practice, acquiring information before bargaining is ubiquitous. For example, homebuyers may rely on inspections and neighborhood research before making an offer, and buyers of artwork or collectibles may privately consult experts about authenticity or condition. Such unverifiable information creates second-order uncertainty for the seller: not only does the buyer not know the seller's type, but the seller also does not know the buyer's information, or equivalently, her belief. As a result, sellers of different types, who hold different beliefs about the buyer's belief, have different incentives, which gives rise to signaling behavior.

In this paper, I study how the buyer's private belief affects efficiency in bargaining over lemons. The seller owns a good that may be of high or low quality and therefore yield either a high or low valuation to the buyer, and this quality is privately and perfectly observed by the seller. The buyer receives a private but imperfect signal about the seller's type and is therefore either optimistic or pessimistic. The seller proposes prices over time, and the buyer decides whether to accept. Since the seller's incentives depend on his belief about the buyer's private type, prices serve not only as potential transaction prices but also as signals of the seller's type. As a result, the seller's offer can affect the buyer's expected valuation of the good.

To examine whether inefficiencies persist as time frictions vanish, in the spirit of the Coase conjecture, I say that trade admits delay in the limit if there exists at least one sequence of equilibria along which delay does not vanish as the discount factor approaches one. I further distinguish between two types of delay. Trade exhibits **weak delay** if there also exists a sequence of equilibria along which delay vanishes, and **strong delay** if no such sequence exists. As shown below, these two forms of delay arise from distinct mechanisms.

As a benchmark, I first show that the model reduces to the classic Coase conjecture outcome when the seller is instead uninformed about the valuation of his good. This is natural, since when the seller cannot signal his type and the buyer's belief is unchanged by the seller's offers, the game reduces to a durable-goods monopoly with privately informed buyers. I then identify two properties that are essential for the Coase conjecture. Property A requires that the on-path market-clearing price be no higher than the off-path one, which naturally holds when the seller has no private information and the buyer's belief never updates following any price history. Property B requires that the seller become increasingly pessimistic about the buyer's type as the latter continues to reject his offers, since only the optimistic buyer would accept a high price. Given the properties described above, if trade were delayed as time frictions vanish, there would always exist a period at which the seller becomes sufficiently pessimistic about the buyer's type that his expected continuation payoff from waiting falls below the market-clearing price, and he therefore prefers to clear the market immediately.

Returning to the main setting, I first investigate the case in which there is no message that reveals the seller's type. The key difference from the benchmark case is that, when the seller can signal his type through prices, a deviation may lead the buyer to believe that the seller is of the low type. I show that this leads to the breakdown of Property A: the on-path market-clearing price, equal to the pessimistic buyer's expected valuation, is now strictly higher than the off-path market-clearing price, equal to the buyer's low-quality valuation, since even the

pessimistic buyer assigns strictly positive probability to the seller being of the high type along the equilibrium path. This gap allows equilibrium beliefs to punish a seller's deviation to clear the market earlier than on the equilibrium path, or equivalently, gives the seller a form of commitment power and thereby sustains delay, which is impossible when the seller is uninformed or the buyer is certain about her valuation. However, this channel is merely permissive: it can sustain delay, but it does not make delay inevitable. Hence, equilibrium delay may be arbitrarily small and is therefore weak by definition.

I then turn to the case in which there exists a message that reveals the seller to be of the low type. I show that trade exhibits strong delay under the Intuitive Criterion. Note that the logic that breaks Property A no longer applies, since the seller must offer the low-quality valuation to clear the market even along the equilibrium path. Instead, the mechanism is closely related to [Feinberg and Skrzypacz \(2005\)](#) (henceforth FS2005). In both settings, the high-type seller knows with certainty that the buyer is optimistic, while the low type cannot rule out a pessimistic buyer. This leads to a breakdown of Property B and allows the high type to separate by delaying offers. Consequently, if delay were to vanish, the high type could separate almost immediately and obtain a payoff arbitrarily close to the high-quality valuation, causing total payoffs to exceed the available surplus. This yields a contradiction, and hence delay is not merely sustainable but uniformly bounded away from zero, making it strong by definition.

A key feature of the present model is that the buyer's expected valuation varies with the seller's price history during the bargaining process, which is characteristic of bargaining over lemons. This is non-trivial and gives rise to a new mechanism generating bargaining inefficiency, namely the breakdown of Property A. It also leads to an intermediate concept between no delay in the limit and persistent delay in the limit, namely weak delay. Together, these results show how the interaction between lemons-market valuation and second-order uncertainty generates forms of delay that are absent in standard Coase environments.

The remainder of the paper is organized as follows. Section 2 discusses the relation of the present paper to the existing literature. Section 3 introduces the model and formally defines the notions of weak and strong delay. Section 4 studies the benchmark case and the two properties underlying the Coase conjecture. Sections 5 and 6 analyze trade delay with and without low-type exclusivity, establish the breakdown of Properties A and B, and derive weak and strong delay, respectively. Section 7 concludes.

2 Related Literature

This work contributes to the literature following [Coase \(1972\)](#), in which trade delay, and thus inefficiency in bargaining, has been a central concern. The key insight of the Coase conjecture, namely immediate trade, is that the seller (or the player who proposes a price) has no commitment power and thus cannot credibly maintain a high price once the buyer rejects past offers when time frictions vanish. A number of papers study how this lack of commitment power can break down, leading to inefficiency even as frictions disappear¹.

¹In most papers, the friction in bargaining comes from the cost of waiting for the next offer, or, more specifically, a non-zero discount rate.

Deneckere and Liang (2006) introduce interdependent values between the seller and the buyer, and show that trade delay can persist in the limit because the proposer’s expected surplus from trade may be negative. Daley and Green (2020) consider bargaining with evolving news, which also gives the price proposer some commitment power to delay trade, since waiting for news is valuable. Board and Pycia (2014) study a setting in which the buyer has an outside option in every period, which bounds the buyer’s utility from below and makes her unwilling to propose a high price immediately. Brzustowski et al. (2023) allow the seller to offer dynamic smart contracts that can retain information reported by the buyer, so abandoning the current contract causes the seller to lose its information content. This gives the seller an incentive to stick to the existing contract and leads to a failure of the Coase conjecture. The present paper instead considers a simpler environment and identifies a new mechanism for the breakdown of the Coase conjecture, namely a gap between the on-path and off-path market-clearing prices.

Other papers focus on the effects of two-sided uncertainty. For example, Cramton (1992) allow the seller to have private costs, making prices a signal of her type. In contrast, the present paper allows the seller’s signaling to affect the buyer’s expected valuation. Abreu and Gul (2000) study reputation effects and allow players to mimic irrational types, making private information part of the strategic posture. In contrast, all strategies in the present paper are endogenous, and delay does not rely on the existence of irrational types. In particular, Feinberg and Skrzypacz (2005) is most closely related to my work. They study a similar setting in which the buyer is privately informed about her valuation, while the seller receives a message about the buyer’s type. Although prices serve as signals, they do not affect the buyer’s valuation. As a result, to sustain persistent delay, their model requires the existence of a conclusive message about a disadvantageous type². When such a message is absent, delay disappears, in contrast to the present paper.

Beyond the bargaining literature, this paper is also related to work on information and inefficiency in the lemons market. In contrast to Akerlof (1978), where asymmetric information leads to inefficiency, Daley and Green (2012) and Daley and Green (2020) show that introducing additional information need not mitigate inefficiency, a conclusion that aligns with the present paper. A common high-level insight in these papers is that delay may arise from information that arrives over time after the trading process has started. By contrast, in my paper all information arrives before bargaining begins, there is no further information arrival during the bargaining process, and the information is private to the buyer rather than public.

3 Model

There is a seller (he) bargaining with a buyer (she). The seller holds a good whose value to the buyer is $v \in \{v_h, v_l\}$, where $v_h > v_l > 0$. The seller’s cost of the good is normalized to zero. The seller is perfectly informed about the quality of the good, while the buyer does not directly observe the valuation of the good and has a non-degenerate prior belief β^{Pr} that the seller is the high type, which is common knowledge.

Before bargaining begins, the buyer receives a private message $m \in \{o, p\}$ about the value (where o stands for optimistic and p for pessimistic). This message determines the buyer’s type. The message structure is given

²The disadvantageous type can be a buyer with a high valuation or a seller with a low-quality good.

by

$$\alpha_0^i \equiv \Pr[m = o \mid i] = 1 - \Pr[m = p \mid i],$$

where $i \in \{h, l\}$ denotes the seller's type. The message is informative if and only if

$$\alpha_0^h > \alpha_0^l,$$

which implies that

$$\alpha_0^h \in (0, 1], \quad \alpha_0^l \in [0, 1).$$

After receiving the message, the buyer updates her belief about the seller's type. The posterior belief upon receiving an optimistic message is

$$\beta_0^o \equiv \frac{\beta^{\text{Pr}} \alpha_0^h}{\beta^{\text{Pr}} \alpha_0^h + (1 - \beta^{\text{Pr}}) \alpha_0^l},$$

while the posterior belief upon receiving a pessimistic message is

$$\beta_0^p \equiv \frac{\beta^{\text{Pr}} (1 - \alpha_0^h)}{\beta^{\text{Pr}} (1 - \alpha_0^h) + (1 - \beta^{\text{Pr}}) (1 - \alpha_0^l)}.$$

It is obvious that

$$0 \leq \beta_0^p < \beta_0^o \leq 1.$$

Throughout the paper, a seller with value v_h is referred to as the **high type**, while one with value v_l is the **low type**. A buyer who receives message o is called the **optimistic type**, and one who receives message p the **pessimistic type**.

Time is discrete and indexed by $t = 0, 1, 2, \dots$. At $t = 0$, the buyer receives the message and forms the corresponding posterior belief about the seller's type. At each period $t \geq 1$, the seller posts a price P , which the buyer can either accept or reject. Both players are risk-neutral expected utility maximizers and share a common discount factor $\delta \in (0, 1)$. If the buyer accepts price P at time t , the seller's payoff is $\delta^t P$ and the buyer's payoff is $\delta^t (v - P)$. If the buyer rejects the offer, bargaining continues in the next period.

Both players are assumed to have perfect recall of the history, that is, the sequence of prices previously offered by the seller and rejected by the buyer:

$$h_\tau = \{P_t\}_{t=1}^\tau.$$

Let $\mathcal{H}$ denote the set of all possible price histories. A buyer of type m chooses the probability of accepting the current offer given the history and the price offer:

$$\sigma^m : \mathcal{H} \times \mathbb{R}^+ \rightarrow [0, 1].$$

A seller of type i chooses a price distribution conditional on the history and the event that the buyer has rejected

all previous offers:

$$\mu^i : \mathcal{H} \rightarrow \Delta(\mathbb{R}^+).$$

Accordingly, the buyer's strategy profile is $\sigma = (\sigma^o, \sigma^p)$ and the seller's strategy profile is $\mu = (\mu^h, \mu^l)$.

The belief of a seller of type i that the buyer is the optimistic type o , conditional on the history, is given by

$$\alpha^i : \mathcal{H} \rightarrow [0, 1].$$

The belief of a buyer of type m that the seller is the high type is given by

$$\beta^m : \mathcal{H} \rightarrow [0, 1].$$

Similarly, denote the belief profile of the players by $\{\alpha, \beta\}$, where $\alpha = (\alpha^h, \alpha^l)$ and $\beta = (\beta^o, \beta^p)$. Furthermore, define the expected value of the good for a buyer given her belief β^m

$$\text{EV}(\beta^m) \equiv \beta^m \cdot v_h + (1 - \beta^m) \cdot v_l.$$

Given the buyer's strategy σ^m and the seller's strategy profile μ , the payoff of a buyer of type m at history h_τ is

$$\begin{aligned} U^m(\mu, \sigma^m | h_\tau) = & \sum_P \left\{ \beta^m(h_\tau) \cdot \mu^h(P | h_\tau) \cdot \left[\sigma^m(h_\tau, P) \cdot (v_h - P) + (1 - \sigma^m(h_\tau, P)) \cdot \delta \cdot U^m(\mu^h, \sigma^m | (h_\tau, P)) \right] \right. \\ & \left. + (1 - \beta^m(h_\tau)) \cdot \mu^l(P | h_\tau) \cdot \left[\sigma^m(h_\tau, P) \cdot (v_l - P) + (1 - \sigma^m(h_\tau, P)) \cdot \delta \cdot U^m(\mu^l, \sigma^m | (h_\tau, P)) \right] \right\}. \end{aligned}$$

Similarly, given the seller strategy μ^i and the buyer's strategy profile σ , the payoff of a seller of type i at history h_τ is

$$\begin{aligned} V^i(\mu^i, \sigma | h_\tau) = & \sum_P \mu^i(P | h_\tau) \cdot \left\{ \alpha^i(h_\tau) \cdot \left[\sigma^o(h_\tau, P) \cdot P + (1 - \sigma^o(h_\tau, P)) \cdot \delta \cdot V^i(\mu^i, \sigma^o | (h_\tau, P)) \right] \right. \\ & \left. + (1 - \alpha^i(h_\tau)) \cdot \left[\sigma^p(h_\tau, P) \cdot P + (1 - \sigma^p(h_\tau, P)) \cdot \delta \cdot V^i(\mu^i, \sigma^p | (h_\tau, P)) \right] \right\}. \end{aligned}$$

I use Perfect Bayesian Equilibrium (PBE) as the solution concept.

Definition 1 (Equilibrium). A quadruple $\mathcal{G} \equiv \{\alpha, \beta, \mu, \sigma\}$ consisting of belief systems $\{\alpha, \beta\}$ and strategy profiles $\{\mu, \sigma\}$ is a Perfect Bayesian Equilibrium if the following conditions hold:

1. **Sequential rationality:** Given beliefs $\{\alpha, \beta\}$, the strategy profiles $\{\mu, \sigma\}$ maximize the expected payoffs of the players. For any $h_\tau \in \mathcal{H}$,

$$\mu^i \in \arg \max_{\mu'} V^i(\mu', \sigma | h_\tau), \quad \sigma^m \in \arg \max_{\sigma'} U^m(\mu, \sigma' | h_\tau).$$

2. **Belief consistency:** Beliefs $\{\alpha, \beta\}$ are updated according to Bayes' rule whenever possible.
3. **Off-path optimality:** Strategies form an equilibrium in every subgame given some beliefs after off-path histories.

To define the delay of trade, let X_S be the indicator of trade. That is, $X_S = 1$ represents that trade occurs at period S and $X_S = 0$ represents the opposite. Therefore given an equilibrium $\mathcal{G}$, the expected total discounted factor until the trade occurs is represented by $\mathbb{E} \left[\delta^{S \cdot 1_{X_S=1}} \mid \mathcal{G} \right]$. I define delay of trade in the following:

Definition 2 (Trade Delay). Let $\{\mathcal{G}_k\}$ be a sequence of equilibria such that (i) for each k , $\mathcal{G}_k$ is an equilibrium with $\delta = \delta_k \in (0, 1)$, and (ii) $\delta_k \rightarrow 1$ as $k \rightarrow \infty$. Trade is said to admit **delay in the limit** if there *exists* a sequence $\{\mathcal{G}_k\}$ and $\varepsilon > 0$ such that

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\delta_k^{S \cdot 1_{X_S=1}} \mid \mathcal{G}_k \right] < 1 - \varepsilon.$$

Trade is said to occur **immediately in the limit** if it never occurs with delay in the limit.

Suppose trade occurs with delay in the limit. Then:

1. Trade exhibits **strong delay** if there exists $\varepsilon > 0$ such that for every sequence $\{\mathcal{G}_k\}$,

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\delta_k^{S \cdot 1_{X_S=1}} \mid \mathcal{G}_k \right] < 1 - \varepsilon.$$

2. Trade exhibits **weak delay** if there exists a sequence $\{\mathcal{G}_k\}$ such that

$$\lim_{k \rightarrow \infty} \mathbb{E} \left[\delta_k^{S \cdot 1_{X_S=1}} \mid \mathcal{G}_k \right] = 1.$$

The analysis proceeds as follows. As a benchmark, I first examine the uninformed-seller case, in which the Coase conjecture naturally holds. I then study trade delay under the assumption that messages are completely mixed across seller types, and show that trade exhibits weak delay. Finally, I consider the case in which the buyer may receive a low-type-exclusive message, similar to FS2005, and show that this leads to strong delay. The contrast between these results highlights the distinct mechanisms generating different types of trade delay.

4 Benchmark: Uninformed Seller and Immediate Trade

I begin with a benchmark case in which the seller does not observe his type, that is, the quality of the product or its valuation to the buyer. Since the seller's strategy cannot depend on his type, the buyer cannot update her belief about the seller's type following any deviation by the seller. As a result, the game exhibits immediate trade in the limit, as in the classical Coase conjecture:

Proposition 1 (Uninformed Seller, No Delay in the Limit). *In the bargaining game where the seller is not informed about his type, trade occurs immediately in the limit and the equilibrium price converges to*

$$P^* = EV(\beta_0^p).$$

Proof. When the seller is uninformed about his type, the game reduces to a durable-goods monopoly with privately informed buyers. By [Coase \(1972\)](#), equilibrium delay vanishes as the discount factor approaches one, and the price converges to $EV(\beta_0^p)$. $\square$

Since this is a classic result, I do not repeat the full intuition here. Instead, I highlight two key properties of the environment, without either of which the Coase conjecture may fail.

Property A. The off-path market-clearing price is no lower than that along the equilibrium path.

Property B. The seller becomes increasingly pessimistic about the buyer's belief as the buyer rejects more offers.

The two properties are straightforward. First, because the buyer never updates her belief in response to the seller's offers, the price $EV(\beta_0^p)$ always completely clears the market, whether on or off the equilibrium path, in the sense that the buyer accepts it for sure regardless of her type and the price history. Second, only the optimistic buyer would accept a price strictly above the pessimistic buyer's expected valuation, so repeated rejections reduce the seller's belief that the buyer has a relatively high expected valuation.

The Coase conjecture relies on these two properties. As the seller becomes increasingly pessimistic, if trade were delayed in the limit, there would exist a period after which the seller assigns an arbitrarily high probability to eventually trading with the pessimistic buyer³. But since the market-clearing price remains unchanged over time, clearing the market immediately is always more profitable than waiting until a later period to do so. Anticipating this, the seller cannot credibly delay trade by maintaining a price above $EV(\beta_0^p)$. Without either of the two properties, delay may remain profitable for the seller, and the Coase conjecture may fail. In the rest of the paper, I show how this logic breaks down and why delay can then persist in the limit.

5 Main Results

In this section, I return to the original setting in which the seller is perfectly informed about his type. I further assume that $\alpha_0^i \in (0, 1)$ for each $i \in \{h, l\}$ ⁴. Under this assumption, the buyer's message space contains no low-type-exclusive message, so even the high-type seller remains uncertain about the buyer's belief. I discuss the opposite case in the next section.

5.1 Equilibrium Construction

I start with introducing a refinement to describe some common properties that fairly restrict the players' beliefs:

Refinement (Condition R). For any offer P at any history h_τ , the buyer's posterior belief profile about the seller's type, denoted by $\beta((h_\tau, P))$, and the seller's posterior belief profile about the buyer's type, denoted by $\alpha((h_\tau, P))$, satisfy

$$\beta^m(h_\tau) = 0 \implies \beta^m((h_\tau, P)) = 0, \quad \alpha^i(h_\tau) = 0 \implies \alpha^i((h_\tau, P)) = 0,$$

³There is also the possibility that the optimistic buyer accepts only with very low probability until the final period, so that the seller remains very optimistic about the buyer's type despite a long history of rejections. However, this also makes waiting strictly suboptimal relative to clearing the market immediately, and therefore cannot occur in equilibrium.

⁴A similar result also holds when $\alpha_0^h \in (0, 1)$ and $\alpha_0^l = 0$, although the proof differs only in inessential ways. I relegate that case to the Appendix A.

and

$$\beta^m(h_\tau) = 1 \implies \beta^m((h_\tau, P)) = 1, \quad \alpha^i(h_\tau) = 1 \implies \alpha^i((h_\tau, P)) = 1,$$

for any $m \in \{o, p\}$ and $i \in \{h, l\}$.

The Condition R requires that if a type is assigned zero probability at some history, then it must continue to receive zero probability at all continuation histories. In other words, the support of the buyer's beliefs at any information set must be contained in the support of beliefs at its preceding information sets. In particular, whenever $\beta^m(h_\tau) = 1$ or $\beta^m(h_\tau) = 0$, the buyer's belief remains unchanged regardless of subsequent events.

I can now establish the following lemma.

Lemma 1. *With condition R, suppose the seller's type is revealed to be l by some price P at history h_τ , trade occurs immediately with probability 1 at price equal to the lower value:*

$$P = v_l,$$

and for $m = o, p$,

$$\sigma^m(h_\tau, P) = 1.$$

Lemma 1 shows that once the seller's type becomes known, the equilibrium outcome coincides with that of a complete-information bargaining game.⁵ This result follows directly from the restriction that the buyer cannot assign positive probability to types that have already been excluded. If $i = l$ becomes common knowledge, trade must occur immediately since any delay would yield the seller a payoff strictly below v_l , which cannot be optimal.

Now consider a class of equilibria, each with a unique on-path price path $\{P_t\}_{t=1}^T$, and such that $\mu^h = \mu^l$; that is,

$$\text{supp } \mu^h(\cdot \mid \{P_t\}_{t=1}^\tau) = \text{supp } \mu^l(\cdot \mid \{P_t\}_{t=1}^\tau) = \{P_{\tau+1}\}$$

for each $\tau = 0, 1, 2, \dots, T-1$. The price sequence is given by

$$P_t = \text{EV}(\beta_0^o) - \delta^{T-t} (\text{EV}(\beta_0^o) - \text{EV}(\beta_0^p)).$$

The buyer's strategy satisfies:

$$\sigma^o(\{P_t\}_{t=1}^\tau, P_{\tau+1}) \in (0, 1), \quad \sigma^p(\{P_t\}_{t=1}^\tau, P_{\tau+1}) = 0 \quad \text{for } \tau \leq T-1,$$

and

$$\sigma^o(\{P_t\}_{t=1}^{T-1}, P_T) = \sigma^p(\{P_t\}_{t=1}^{T-1}, P_T) = 1.$$

⁵In a bargaining game with complete information, the unique equilibrium outcome is that the seller offers a price equal to the value of the good and the buyer accepts immediately. To see this, note that the seller can credibly stick to a price arbitrarily close to but strictly lower than the buyer's valuation, which the buyer would accept immediately. Thus the equilibrium leaves the buyer zero surplus.

The buyer's belief is constant because there is no separation of seller, and therefore for any on-path history h_τ ,

$$\beta^m(h_\tau) = \beta^m(h_0) = \beta_0^m.$$

If the seller deviates from the prescribed price path, the buyer assigns probability one to the seller being the low type and therefore only accepts price weakly lower than v_l .

The seller's belief is pinned down by

$$\alpha^i(h_\tau = \{P_t\}_{t=1}^\tau) = \frac{\alpha_0^i \cdot \prod_{t'=1}^\tau (1 - \sigma^o(\{P_t\}_{t=1}^{t'-1}, P_{t'}))}{\alpha_0^i \cdot \prod_{t'=1}^\tau (1 - \sigma^o(\{P_t\}_{t=1}^{t'-1}, P_{t'})) + 1 - \alpha_0^i}$$

for any on-path history h_τ . Note that in this equilibrium the only seller's off-path belief need to specify is after the buyer reject the offer P_T after $\{P_t\}_{t=1}^{T-1}$. In this case, the seller always believes that the buyer is the pessimistic type.

Lastly, I further impose that

$$1 < T \leq \left\lfloor \log_\delta \frac{v_l}{P_T} \right\rfloor = \left\lfloor \log_\delta \frac{v_l}{\text{EV}(\beta_0^p)} \right\rfloor. \quad (1)$$

This implies that for any history h_τ , the i -type seller's continuation payoff satisfies

$$\begin{aligned} V^i(\mu^i, \sigma | h_\tau) &= \alpha^i(h_\tau) \cdot \sigma^o(h_\tau, P_{\tau+1}) \cdot P_{\tau+1} + (1 - \alpha^i(h_\tau) \cdot \sigma^o(h_\tau, P_{\tau+1})) \cdot \delta \cdot V^i(\mu^i, \sigma | (h_\tau, P_{\tau+1})) \\ &\geq \delta \cdot V^i(\mu^i, \sigma | (h_\tau, P_{\tau+1})) \geq \dots \geq \delta^{T-\tau} P_T \geq \delta^T P_T \geq v_l. \end{aligned}$$

By Lemma 1, if the seller deviates, he always get payoff equal to v_l , which is always lower than the equilibrium payoff. Thus no seller type has a profitable deviation. In Appendix B, I also provide a robustness check to verify that the constructed equilibrium survives the D1 Criterion.

This equilibrium is uniquely characterized by the price path $\{P_t\}_{t=1}^T$ and the sequence of acceptance probabilities of the optimistic buyer, $\sigma^o \in (0, 1)^{T-1}$. I denote it by $\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}$. Since the trade happens immediately after the buyer's type is revealed to be low, it is easy to see that this equilibrium satisfies the condition R.

5.2 Weak Delay under No Low-Type Exclusivity

I now formally show that the absence of a low-type-exclusive message for the buyer leads to weak delay:

Theorem 1 (Weak Delay without Low-Type Exclusivity). *Under Condition R, if $\alpha_0^i \in (0, 1)$ for each $i \in \{h, l\}$, then trade exhibits **weak delay**.*

Proof. By the definition of weak delay in Definition 2, it is sufficient to establish two arguments.

- (i) Delay is bounded away from zero along some equilibrium sequence. That is, there exists a sequence of equilibria $\{\mathcal{G}_k\}$ such that the delay of trade is strictly bounded away from zero. To this end, for each

$\delta \in (0, 1)$ and equilibrium $\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}$, let

$$T = \left\lfloor \log_{\delta} \frac{v_l}{\text{EV}(\beta_0^p)} \right\rfloor > 1.$$

Note that this satisfies (1) and is always feasible for δ sufficiently closed to 1, given that $\text{EV}(\beta_0^p) > v_l$ and therefore $\lim_{\delta \rightarrow 1} \log_{\delta} \frac{v_l}{\text{EV}(\beta_0^p)} = \infty$.

To bound the expected discount factor at the time of trade from above, let $\sigma^o(h_0, P_1) \rightarrow 1$. Then

$$\begin{aligned} \lim_{\delta \rightarrow 1} \mathbb{E} \left[\delta^{S \cdot 1_{X_S=1}} \mid \mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o} \right] &\leq \lim_{\delta \rightarrow 1} \left\{ \left[\beta^{\text{Pr}} \cdot \alpha_0^h + (1 - \beta^{\text{Pr}}) \cdot \alpha_0^l \right] \cdot 1 + \left[\beta^{\text{Pr}} \cdot (1 - \alpha_0^h) + (1 - \beta^{\text{Pr}}) \cdot (1 - \alpha_0^l) \right] \cdot \delta^T \right\} \\ &= \left[\beta^{\text{Pr}} \cdot \alpha_0^h + (1 - \beta^{\text{Pr}}) \cdot \alpha_0^l \right] + \left[\beta^{\text{Pr}} \cdot (1 - \alpha_0^h) + (1 - \beta^{\text{Pr}}) \cdot (1 - \alpha_0^l) \right] \cdot \frac{v_l}{\text{EV}(\beta_0^p)} \\ &= 1 - \left[\beta^{\text{Pr}} \cdot (1 - \alpha_0^h) + (1 - \beta^{\text{Pr}}) \cdot (1 - \alpha_0^l) \right] \cdot \left(1 - \frac{v_l}{\text{EV}(\beta_0^p)} \right) < 1. \end{aligned}$$

- (ii) Delay converges to zero along some equilibrium sequence. That is, there exists a sequence of equilibria $\{\mathcal{G}_k\}$ such that the delay of trade converges to zero. To this end, again for each $\delta \in (0, 1)$ and equilibrium $\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}$, fix $T = 2$. Note that we can get the lower bound of expected discount factor by letting $\sigma^o(h_0, P_1) \rightarrow 0$:

$$\begin{aligned} \lim_{\delta \rightarrow 1} \mathbb{E} \left[\delta^{S \cdot 1_{X_S=1}} \mid \mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o} \right] &\geq \lim_{\delta \rightarrow 1} \left\{ \left[\beta^{\text{Pr}} \cdot \alpha_0^h + (1 - \beta^{\text{Pr}}) \cdot \alpha_0^l \right] \cdot \delta + \left[\beta^{\text{Pr}} \cdot (1 - \alpha_0^h) + (1 - \beta^{\text{Pr}}) \cdot (1 - \alpha_0^l) \right] \cdot \delta^2 \right\} \\ &= 1. \end{aligned}$$

Combining (i) and (ii), trade exhibits weak delay. □

The proof of Theorem 1 highlights the distinction between weak delay and the absence of delay in the limit. In particular, weak delay is characterized by the existence of a sequence of equilibria along which trade delay does not vanish, as well as the existence of another sequence along which it vanishes. In the following, I show that this contrasts sharply with the benchmark case with uninformed seller.

The key difference between this model and the benchmark Coase environment lies in the failure of **Property A** discussed in Section 4. In standard Coase-style models, the seller cannot influence the buyer's expected valuation through his actions. This may be because the seller lacks private information, as in the benchmark model, or because the buyer's (expected) valuation is fixed independently of the seller's behavior, as in Coase (1972). As a result, the seller can clear the market at any time using the same market-clearing price.

In contrast, when the seller is informed about his type, the buyer's beliefs following a deviation can be strictly more pessimistic than those on the equilibrium path, thereby inducing lower expected valuations, potentially as low as v_l , whereas along the equilibrium path the seller can still secure $\text{EV}(\beta_0^p)$ in the future. This creates a wedge between on-path and off-path incentives: even a pessimistic seller may prefer to remain on the equilibrium path rather than deviate to accelerate trade, despite the discounting cost of delay. Consequently, delay can be sustained in equilibrium.

However, the breakdown of Property A alone does not make delay inevitable. It merely permits sufficiently severe off-path punishments that can deter the seller from accelerating trade when equilibrium delay persists in the limit, but it does not rule out arbitrarily small delay or even immediate trade. As a result, it generates weak delay in a **permissive** sense: delay can be sustained, but need not be bounded away from zero. In the next section, I show how the breakdown of Property B instead leads to persistent delay in a stronger channel.

6 Strong Delay and Comparison with FS2005

Now I consider the case in which the buyer receives a conclusive message indicating that the seller is of the low type, that is, $\alpha_0^h = 1$, $\alpha_0^l \in (0, 1)$ and $\beta_0^p = 0$, $\beta_0^o \in (0, 1)$. In other words, the pessimistic buyer is surely that the seller is low type, while the high type seller believes that the buyer should be the optimistic one with probability 1. It is obvious that now the breakdown of Property A we have in Section 5 is no longer valid. However, in the following I show that, now the high type seller has a much stronger incentive to wait for long enough time to signal his type given there is no "private belief" of buyer from his perspective, leading to a breakdown of **Property B**.

Theorem 2 (Strong Delay with Low-Type Exclusivity, Theorem 1 in FS2005). *Suppose there exists a message revealing that the seller is the low type, i.e., $\alpha_0^l \in (0, 1)$ and $\alpha_0^h = 1$. Then trade occurs with **strong delay** under Condition R and the Intuitive Criterion.*

Proof. Let θ satisfy $\delta^\theta = 1 - \frac{1}{n}$ for positive integer n . Suppose not, that is, for any $\delta < 1$, there exists a corresponding equilibrium and n such that

$$\mathbb{E} \left[\delta^{S \cdot \mathbf{1}_{X_S=1}} \right] \geq 1 - \frac{1}{n^2}.$$

By definition, we have

$$\mathbb{E} \left[\delta^{S \cdot \mathbf{1}_{X_S=1}} \right] \leq \left(1 - \frac{1}{n} \right) \cdot \text{Prob.} \left[\delta^{S \cdot \mathbf{1}_{X_S=1}} \leq \delta^\theta \right] + \left(1 - \Pr \left[\delta^{S \cdot \mathbf{1}_{X_S=1}} \leq \delta^\theta \right] \right).$$

Combine the two inequalities above, we have

$$\text{Prob.} \left[\delta^{S \cdot \mathbf{1}_{X_S=1}} \leq \delta^\theta \right] = \text{Prob.} \left[\theta \leq S \cdot \mathbf{1}_{X_S=1} \right] \leq \frac{1}{n},$$

that is, the probability that the game does not terminate by period θ is weakly lower than $\frac{1}{n}$. Now, consider the payoff of the low type seller if she mimics the high type one for θ periods, and note that the low type's equilibrium payoff should be weakly larger than it, we have

$$V^l(\mu^l, \sigma) \geq \alpha_0^l \left(1 - \frac{1}{n} \right) \left(V^h(\mu^h, \sigma) - \frac{v_h}{n} \right) + (1 - \alpha_0^l) \left(1 - \frac{1}{n} \right) v_l,$$

which consists of two components: one is the expected payoff of the high type seller by period θ conditional on

that the buyer is optimistic and weakly larger than $V^h(\mu^h, \sigma) - \frac{v_h}{n}$ because the high type seller at most get v_h with probability $1/n$ after period θ ; the other is the expected payoff when the buyer is pessimistic, which is at least v_l . Now pickup a period $\tilde{T}$, such that

$$\begin{aligned} \delta^{\tilde{T}+1}(\alpha_0^l v_h + (1 - \alpha_0^l) v_l) &< \alpha_0^l \left(1 - \frac{1}{n}\right) \left(V^h(\mu^h, \sigma) - \frac{v_h}{n}\right) + (1 - \alpha_0^l) \left(1 - \frac{1}{n}\right) v_l, \\ &\leq \delta^{\tilde{T}}(\alpha_0^l v_h + (1 - \alpha_0^l) v_l) \end{aligned} \quad (2)$$

which is always possible given $V^h(\mu^h, \sigma) \leq v_h$. Now consider a deviation for the high type seller: she keeps making an unserious offer for $\tilde{T}$ period. Since for the low type seller, the best she can get after $\tilde{T}$ periods' unserious offer is exactly $\alpha_0^l v_h + (1 - \alpha_0^l) v_l$, this deviation is never profitable and reveals the seller's type to be high type by Intuitive Criterion. Thus, we may conclude that

$$V^h(\mu^h, \sigma) \geq \delta^{\tilde{T}+1} v_h.$$

Combine the two inequalities above, we have

$$V^h(\mu^h, \sigma) \geq \delta v_h \frac{\alpha_0^l \left(1 - \frac{1}{n}\right) \left(-\frac{v_h}{n}\right) + (1 - \alpha_0^l) \left(1 - \frac{1}{n}\right) v_l}{(\alpha_0^l v_h + (1 - \alpha_0^l) v_l) - \delta v_h \alpha_0^l \left(1 - \frac{1}{n}\right)} \xrightarrow{\delta \rightarrow 1, n \rightarrow \infty} v_h.$$

Now, note that since the buyer's expected payoff needs to be weakly larger than 0, we must have

$$\beta^{\text{Pr}} \cdot V^h(\mu^h, \sigma) + (1 - \beta^{\text{Pr}}) \cdot V^l(\mu^l, \sigma) \leq \beta^{\text{Pr}} v_h + (1 - \beta^{\text{Pr}}) v_l.$$

By letting $n \rightarrow \infty$, we have

$$\beta^{\text{Pr}} v_h + (1 - \beta^{\text{Pr}})(\alpha_0^l v_h + (1 - \alpha_0^l) v_l) \leq \beta^{\text{Pr}} v_h + (1 - \beta^{\text{Pr}}) v_l,$$

or,

$$\alpha_0^l (1 - \beta^{\text{Pr}})(v_h - v_l) \leq 0,$$

a contradiction to $\alpha_0^l > 0$, $\beta^{\text{Pr}} < 1$ and $v_h - v_l > 0$. □

When there exists a message that excludes the high type, trade exhibits strong delay, in the sense that delay is strictly bounded away from zero along any sequence of equilibria. The key driving force behind the result lies in the failure of Property B: even after any history of rejections, the high-type seller knows with certainty that the buyer is optimistic and that her belief remains flexible. In contrast, the low-type seller must always assign positive probability to a pessimistic buyer whose belief is stubborn and never updated. This asymmetry allows the high type to separate through delay once the Intuitive Criterion is imposed, because the buyer must then assign belief to the only type that can benefit from such delay, namely the high type. Since the breakdown of Property B effectively restricts the buyer's belief following unexpected delay when the market is quickly cleared

on the equilibrium path, it gives the seller an incentive to signal his type through delay and therefore generates strong delay in a **forcing** sense. This contrasts with the weak delay generated by the failure of Property A.

The mechanism described above is similar in spirit to FS2005, where the buyer knows her valuation while the seller receives a private signal about the buyer's valuation, so that high- and low-type sellers differ in how optimistic they are about the buyer's persistent valuation. There is, however, a difference in why the high-type seller can obtain v_h once he reveals his identity. In the present paper, this is natural because the buyer's expected valuation depends on her belief about the seller's type. In FS2005, by contrast, the buyer's expected valuation is independent of the seller's type. Instead, once it becomes common knowledge that the seller is of the high type, the high-valuation buyer no longer has an incentive to withhold trade in the hope that the seller is unaware of her type.

When there is no low-type exclusivity, as in Section 5, the above logic disappears because the high-type seller is no longer certain about the buyer's belief, and the pessimistic buyer no longer assigns probability 1 to the seller being of the low type. As a result, if the high type can convince the buyer that he is of the high type, then the low type can do so as well⁶, so the high type no longer has a valid incentive to separate by signaling his type. Instead, the possibility of persistent delay relies on the failure of Property A. This contrasts with FS2005, where the absence of low-type exclusivity eliminates the possibility of non-vanishing trade delay. The key difference lies in the structure of the buyer's valuation. In the present paper, valuation uncertainty is in the spirit of the lemons market: the buyer's expected valuation depends on her belief about the seller's type. In FS2005, by contrast, uncertainty comes from heterogeneity in the buyer's preference, and the seller receives a private message about the buyer's valuation. As a result, although prices still serve as signals of the seller's type, they no longer affect the buyer's expected valuation, so Property A remains valid, yielding an outcome that resembles the classic Coase result.

7 Conclusion

This paper identifies two properties that jointly deliver the Coase conjecture, and shows how each can fail in an environment with lemons-market valuation and second-order uncertainty. While one of these channels has been studied in the existing literature, the other contributes a new theoretical insight: even completely mixed information collected by the buyer may generate trade delay. The key force is that the buyer's expected valuation depends on her belief about the seller's type, and this belief can in turn be influenced by the seller's pricing behavior. As a result, the off-path market-clearing price can be lower than the on-path one, which gives the seller a form of commitment power and sustains delay. This perspective is also consistent with practice, where the information privately acquired by buyers is often imperfect rather than conclusive, and yet significant trade delay may still arise.

The distinction between the failures of different properties also gives rise to different forms of trade delay,

⁶This modifies (2) in the proof of Theorem 2. Without low-type exclusivity, the low-type seller can also guarantee v_h after withholding offers for $\tilde{T}$ periods. Hence the lower and upper bounds in (2) should be replaced by $\delta^{\tilde{T}+1}v_h$ and $\delta^{\tilde{T}}v_h$, respectively.

namely weak delay and strong delay, which further clarify the efficiency consequences of second-order uncertainty under different information structures. Strong delay caused by low-type exclusivity implies a more robust welfare loss, since delay persists along every equilibrium sequence as frictions vanish. Weak delay under completely mixed messages, however, is also practically relevant, because real bargaining environments often admit multiple plausible patterns of behavior under the same fundamentals. In this sense, weak delay captures a more fragile, but still realistic, form of inefficiency that depends on how market participants coordinate their beliefs and strategies.

References

- Abreu, Dilip and Faruk Gul (2000) "Bargaining and reputation," *Econometrica*, 68 (1), 85–117.
- Akerlof, George A (1978) "The market for "lemons": Quality uncertainty and the market mechanism," in *Uncertainty in economics*, 235–251: Elsevier.
- Banks, Jeffrey S and Joel Sobel (1987) "Equilibrium selection in signaling games," *Econometrica: Journal of the Econometric Society*, 647–661.
- Board, Simon and Marek Pycia (2014) "Outside options and the failure of the Coase conjecture," *American Economic Review*, 104 (2), 656–71.
- Brzustowski, Thomas, Alkis Georgiadis-Harris, and Balázs Szentes (2023) "Smart contracts and the coase conjecture," *American Economic Review*, 113 (5), 1334–1359.
- Coase, Ronald H (1972) "Durability and monopoly," *The Journal of Law and Economics*, 15 (1), 143–149.
- Cramton, Peter C (1992) "Strategic delay in bargaining with two-sided uncertainty," *The Review of Economic Studies*, 59 (1), 205–225.
- Daley, Brendan and Brett Green (2012) "Waiting for News in the Market for Lemons," *Econometrica*, 80 (4), 1433–1504.
- (2020) "Bargaining and news," *American Economic Review*, 110 (2), 428–74.
- Deneckere, Raymond and Meng-Yu Liang (2006) "Bargaining with interdependent values," *Econometrica*, 74 (5), 1309–1364.
- Feinberg, Yossi and Andrzej Skrzypacz (2005) "Uncertainty about uncertainty and delay in bargaining," *Econometrica*, 73 (1), 69–91.

A High-Type Exclusivity

In this section, I study the omitted environment in which $\alpha_0^h \in (0, 1)$ and $\alpha_0^l = 0$. In this case, the optimistic buyer knows with certainty that the seller is of the high type, while the pessimistic buyer remains uncertain; equivalently, $\beta_0^o = 1$ and $\beta_0^p \in (0, 1)$. I show that this case yields a result similar to the completely mixed-message case studied in Section 5, namely:

Proposition 2 (Weak Delay with High-Type Exclusivity). *Under Condition R, if $\alpha_0^h \in (0, 1)$ and $\alpha_0^l = 0$, then trade exhibits *weak delay*.*

Proof. The only difference from the proof of Theorem 1 is that, following a deviation by the seller, Condition R allows only the pessimistic buyer to believe that the seller is of the low type, since the optimistic buyer already knows with certainty that the seller is of the high type. In other words, the continuation game after the seller's deviation becomes payoff-equivalent to the one studied in FS2005, where the buyer privately knows her valuation while the seller holds private beliefs about the buyer's type but is never certain that the buyer has a high valuation.

By Theorem 2 in FS2005, as $\delta \rightarrow 1$, trade delay vanishes in this continuation game and the seller's payoff converges to v_l . Therefore, for any given beliefs of the players α, β and any constant $\varepsilon > 0$, there exists some $\underline{\delta} \in (0, 1)$ such that whenever $\delta \in (\underline{\delta}, 1)$, if the seller deviates to a price $P'_{\tau+1}$ at history h_τ , then

$$V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1})) - v_l < \varepsilon,$$

where $V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1}))$ denotes the seller's off-path payoff following the deviation.

Hence, for sufficiently large δ , we can still construct an equilibrium of the type described in the proof of Theorem 1, with

$$V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1})) = v_l + \varepsilon < \text{EV}(\beta_0^p)$$

for any on-path history h_τ and off-path offer $P'_{\tau+1}$. It follows that

$$T = \left\lceil \log_\delta \frac{V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1}))}{\text{EV}(\beta_0^p)} \right\rceil > 1$$

exists. The remainder of the proof then follows exactly as in the proof of the main result. $\square$

B D1 Criterion

In this section, I impose the D1 Criterion introduced by Banks and Sobel (1987) to examine the robustness of the equilibrium construction. This is important because the result of persistent trade delay in the limit depends on the existence of such equilibria. I use the D1 Criterion as a counterpart to the Intuitive Criterion imposed in FS2005, since the latter imposes only limited restrictions in the case without low-type exclusivity. The D1 Criterion plays a central role in the literature, as it rules out equilibria supported by deviations that are more

attractive to the high-type seller than to the low type. Intuitively, if a deviation yields strictly higher gains for the high type than for the low type, then the buyer assigns probability one to the seller being the high type after observing that deviation, so such an equilibrium cannot be sustained.

To apply D1 in a dynamic setting, we require a dynamic analogue of the static criterion. In contrast to the static case, where off-path beliefs are evaluated against a set of mixed responses, here the continuation game following a deviation must itself constitute a Perfect Bayesian Equilibrium (PBE). Thus, the relevant object is the set of continuation PBEs, which differ only in off-path beliefs.

Fix an equilibrium $\mathcal{G}$, an on-path history h_τ , and a deviation P' . Let $\Sigma_i^0(\mathcal{G}, h_\tau, P')$ and $\Sigma_i(\mathcal{G}, h_\tau, P')$ denote, respectively, the sets of buyer strategies in continuation PBEs such that type $i \in \{h, l\}$ seller is weakly or strictly better off by deviating to P' than by staying on the equilibrium path. The D1 Criterion requires that if

$$\Sigma_l^0(\mathcal{G}, h_\tau, P') \subset \Sigma_h(\mathcal{G}, h_\tau, P'),$$

then the buyer assigns probability one to the seller being the high type after observing P' , and the equilibrium $\mathcal{G}$ is eliminated.

I now show that, the equilibrium I constructed in the last subsection can further survive the D1 criterion.

Proposition 3. *The equilibrium $\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^0}$ survives the D1 Criterion.*

Proof. First consider the terminal history $\tau = T$. At this point, the buyer's type is effectively revealed to be pessimistic, so the continuation game is one of complete information. Both seller types obtain the same payoff, and thus no deviation can be supported by beliefs that would eliminate the equilibrium.

Now fix any $\tau < T$. To verify D1, it suffices to show that there exists an off-path continuation equilibrium under which the low-type seller weakly benefits from a deviation while the high-type seller does not. To this end, consider a deviation to some price $P'_{\tau+1}$ and construct an off-path equilibrium

$$\mathcal{G}' = \mathcal{G}_{\{P'_t\}_{t=1}^T, \sigma'^0}.$$

We choose the off-path beliefs β' and horizon T' such that:

- (i) The continuation game after $(h_\tau, P'_{\tau+1})$ follows $\mathcal{G}'$;
- (ii) The low-type seller is indifferent between deviating and staying on-path:

$$V^l(\mu^{ll}, \sigma' | (h_\tau, P'_{\tau+1})) = V^l(\mu^l, \sigma | (h_\tau, P_{\tau+1}));$$

- (iii) The high-type seller strictly prefers not to deviate:

$$V^h(\mu^{hh}, \sigma' | (h_\tau, P'_{\tau+1})) < V^h(\mu^h, \sigma | (h_\tau, P_{\tau+1})).$$

We now argue that such a construction is always feasible. Since the continuation equilibrium is off-path, the

buyer's beliefs β' can be chosen freely subject to the condition R. In particular, the off-path beliefs of the buyer should satisfy

$$\beta^{o'}(h_\tau, P'_{\tau+1}) > \beta^{p'}(h_\tau, P'_{\tau+1}).$$

Now, if setting $\sigma^{o'}(h_\tau, \cdot)$ sufficiently close to 0 along the path for $\tau \leq T' - 1$, so that $V^h(\mu^{h'}, \sigma' \mid (h_\tau, P'_{\tau+1})) \rightarrow \delta^{T'-\tau-1} \cdot P'_{T'}$ by the settings of equilibrium. By doing so, we can make the difference between the high-type seller's payoff and that of the low type arbitrarily close to 0. That is, for any $\varepsilon > 0$, it is always possible to have

$$V^h(\mu^{h'}, \sigma' \mid (h_\tau, P'_{\tau+1})) - V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1})) \leq \varepsilon.$$

Furthermore, note we can make the continuation payoff of the low-type seller arbitrarily close to zero by setting $T' \rightarrow \infty$, and arbitrarily close to v_h by setting $\beta^{p'}(h_\tau, P'_{\tau+1}) \rightarrow 1$ and $T' \rightarrow \tau$. Then, by continuity, there exists $\tau < T' < \infty$ such that the low-type seller is exactly indifferent, or

$$V^l(\mu^l, \sigma \mid (h_\tau, P_{\tau+1})) = V^l(\mu^{l'}, \sigma' \mid (h_\tau, P'_{\tau+1})).$$

Note that this also implies that both seller types has no incentive to deviate to offer v_l . Now, since there exists $\varepsilon > 0$ such that

$$V^h(\mu^h, \sigma \mid (h_\tau, P_{\tau+1})) - V^l(\mu^l, \sigma \mid (h_\tau, P_{\tau+1})) > \varepsilon$$

by the definition of $\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}$, the high type's expected payoff under $\mathcal{G}'$ is strictly lower than that on the equilibrium path; that is,

$$\sigma' \notin \Sigma_h(\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}, h_\tau, P'),$$

but

$$\sigma' \in \Sigma_l^0(\mathcal{G}_{\{P_t\}_{t=1}^T, \sigma^o}, h_\tau, P').$$

Therefore, the deviation is not necessary more attractive for the high type than for the low type. This violates the D1 condition, and thus the equilibrium survives D1. $\square$