ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute Discussion Paper No. 284

Bargaining with Binary Private Information

Francesc Dilmé

March 2024

www.econtribute.de

UNIVERSITÄT BONN
Universität zu Köln

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866 is gratefully acknowledged.

Bargaining with Binary Private Information

Francesc Dilmé*

5th March 2024

This paper studies bargaining between a seller and a buyer with binary private valuation. Because the setting is more tractable than the case of general valuation distributions (studied in Gul et al., 1986), we are able to explicitly construct the full set of equilibria via induction. This lets us provide a simple proof of the Coase conjecture and obtain new results: The seller extracts all surplus as she becomes more patient, and the equilibrium outcome converges to the perfect-information outcome as private information vanishes. We also fully characterize the case where there is a deadline: We establish that if the probability that the buyer's valuation is high is large enough, then the seller charges a high price at all times, there are trade bursts at the outset and the deadline, and trade occurs at a constant rate in between.

Keywords: Bargaining, private information, one-sided offers.

JEL Classifications: C78, D82

*University of Bonn; fdilme@uni-bonn.de. This work was funded by a grant from the European Research Council (ERC 949465). Support from the German Research Foundation (DFG) through CRC TR 224 (Project B02) and under Germany's Excellence Strategy – EXC 2126/1–390838866 and EXC 2047–390685813 – is gratefully acknowledged.

1 Introduction

In a seminal paper, Gul et al. (1986) studied the problem faced by a monopolist selling durable goods over time to a mass of buyers with different valuations. They showed that the Coase conjecture (Coase, 1972) holds when there is a gap between the seller's cost and the support of the buyers' valuations: As the frequency with which the monopolist can change prices increases, all equilibrium prices converge to the lowest buyer valuation. The same result applies in a setting where a seller and a single buyer bargain over the price of a durable good, the buyer has private information on his valuation, and the seller makes price offers over time.

An important part of the subsequent bargaining literature has focused on cases where demand is binary. Binary demand has been extensively used in the study of dynamic monopolistic markets with arriving buyers (Sobel, 1991), bargaining in decentralized markets (Lauermann and Wolinsky, 2016), bargaining with arriving buyers (Kaya and Kim, 2018), revenue management (Dilmé and Li, 2019), bargaining with news (Daley and Green, 2020), bargaining with divisible goods (Gerardi et al., 2022), and repeated bargaining (Kaya and Roy, 2022, and Dilmé, 2023a). The typical reasons for focusing on binary demand are simplicity (it provides a minimal, canonical departure from perfect information), solvability (it is sometimes necessary to make the model tractable), and clarity (it allows for more straightforward arguments and closed-form solutions).

This paper studies a bargaining setting in which a seller makes price offers to a buyer with a private binary valuation. Even though the analysis in Gul et al. (1986) for a general valuation distribution accommodates the case of a binary distribution, we believe it is important to study the binary setting in isolation, for several reasons.

First, in the binary setting we can explicitly and completely characterize the equilibrium behavior while avoiding most of the technicalities required in the general case, which often cloud the exposition. Although our proof is not short, it is intuitive and self-contained. The key step is to pin down the maximal range of the seller's prior beliefs about the buyer's valuation such that, in all equilibria, the seller offers a low price in the first period and the buyer accepts it immediately. From this first range, one can inductively identify subsequent ranges of priors such that, in all equilibria, whenever the prior lies within the given range, the seller offers a given price for sure in the first period. The Coase conjecture then follows from a simple argument using the closed-form expressions describing the equilibrium outcome.

Second, the tractability of the setting permits us to study the case where the seller's and buyer's discount rates differ, for which we obtain new comparative-statics results. We show that, in the limit as the seller becomes more patient, the Coasian forces weaken, letting the seller extract the full trade surplus from the buyer. Similarly, as the prior degenerates towards a high valuation, the equilibrium outcome coincides with the unique outcome of the perfect-information game, so the buyer fails to benefit from reputation effects. We also obtain the seller's optimal pricing under commitment, which is generically deterministic and features price discrimination when the seller is more patient than the buyer.

Third, our analysis lets us obtain the equilibrium dynamics in the case where the horizon is finite; hence we are able to compare the infinite- and finite-horizon cases within the same framework. We completely characterize the equilibrium behavior in the finite-horizon case, now using backward induction from the deadline. We show that, as the frequency with which offers are made increases, the equilibrium outcome is determined by a threshold in the seller's prior belief about the buyer's having the high valuation. If the seller's prior is lower than the threshold, the Coasian outcome – in which trade occurs immediately at a price equal to the low valuation – is the unique equilibrium outcome. On the other hand, if the seller's prior is higher than the threshold, the unique equilibrium outcome is drastically different: Unlike in the no-gap case (see Fuchs and Skrzypacz, 2013, and Dilmé, 2023b), there is a burst of trade at the outset at a price equal to the high valuation. Then, at all times before the deadline, the price equals the high valuation, and trade occurs at a constant rate. Finally, there is another trade burst at the deadline, also at a price equal to the high valuation.

Overall, we provide a complete analysis of monopoly pricing with binary demand, and we obtain several new results. Our work sheds light on a phenomenon that is central to a large body of literature. We hope our analysis can be extended beyond the classical setting, for example, in bargaining with news or endogenous types.

Structure of the paper: In Section 2 we present the model. In Sections 3 and 4 we analyze the infinite-horizon and finite-horizon cases, respectively. Section 5 concludes. The appendix contains the proofs of all the results.

2 Model

A seller (she) and a buyer (he) bargain over time over the price of an indivisible durable good.[1] The buyer's private valuation for the good is θ, which is either high (h), with probability $\phi_0 \in (0,1)$, or low (ℓ), with probability $1-\phi_0$, with $h > \ell > 0$.

Time is discrete, $t \in \mathcal{T} := \{0,1,...,T\}$, where either $T = +\infty$ (infinite horizon) or $T \in \mathbb{Z}_+$ (finite horizon). In each time period, the seller offers a price $p \in [0,h]$.[2] The buyer either accepts the offer, in which case the game ends, or rejects it, in which case the game continues. The discount factors of the seller and the buyer are $\delta_s \equiv e^{-r_s \Delta}$ and $\delta_b \equiv e^{-r_b \Delta}$, respectively; we interpret $r_s, r_b > 0$ as discount rates and $\Delta > 0$ as the length of all periods. If the buyer accepts price p in period t, he gets $\delta_b^t (\theta - p)$, and the seller gets $\delta_s^t p$. If the buyer never accepts an offer, then both the buyer and the seller get 0. If trade occurs in period t, we will sometimes say it occurs at "physical time $t\Delta$."

2.1 Strategies and equilibrium concept

For each $t \in \mathcal{T}$, a *t-history* of the game is a finite sequence of prices $p^t \equiv (p_0, ..., p_{t-1}) \in [0,h]^t$. We let $H := \cup_{t=0}^{T} [0,h]^t$ be the set of histories. A *strategy of the seller* is a map taking each history to a distribution over price offers, $\pi : H \to \Delta([0,h])$. For each $\theta \in \{\ell, h\}$, the *strategy of the θ-buyer* is a map taking each combination of history and price offer to a probability of accepting the offer, $\alpha_\theta : H \times [0,h] \to [0,1]$.

A *belief system* is a map $\phi : H \to [0,1]$, where $\phi(p^t)$ is interpreted as the probability the seller assigns to the buyer's valuation being h at history p^t. An *assessment* is a pair composed of a strategy profile and a belief system. Given an assessment consisting of a strategy profile $(\pi, \alpha_\ell, \alpha_h)$ and a belief system ϕ, the continuation payoffs of the seller and the buyer, respectively, after the history p^t are given by

$$C(p^t; \pi, \alpha_\ell, \alpha_h, \phi) := \phi(p^t)\, \mathbb{E}\left[\delta_s^{\tilde{\tau}-t}\, \tilde{p}_{\tilde{\tau}} \middle| p^t; \pi, \alpha_h\right] + (1-\phi(p^t))\, \mathbb{E}\left[\delta_s^{\tilde{\tau}-t}\, \tilde{p}_{\tilde{\tau}} \middle| p^t; \pi, \alpha_\ell\right]$$

[1] Unlike Gul et al. (1986), and following the convention in most of the bargaining literature, we analyze a bargaining setting with a seller and a privately informed buyer. As pointed out in Ausubel and Deneckere (1989), the model is mathematically equivalent to one where a monopolist sells to an atomless market.

[2] The assumption that the set of possible prices is bounded is necessary to guarantee a well-defined payoff for any strategy profile. Moreover, even if the set of prices were larger, only prices in the range $[0,h]$ would be relevant for studying equilibrium behavior.

and, for each $\theta \in \{\ell, h\}$,

$$V_\theta(p^t;\pi,\alpha_\theta) := \mathbb{E}\left[\delta_s^{\tilde{\tau}-t}\,(\theta - \tilde{p}_{\tilde{\tau}})\,\middle|\,p^t;\pi,\alpha_\theta\right]\ ;$$

here $\tilde{\tau}$ is the transaction time, and $\mathbb{E}$ is the expectation with respect to $\tilde{\tau}$, whose distribution is fully determined by the assessment and history. We will look for *perfect Bayesian equilibria*, henceforth referred to simply as *equilibria*.

Definition 2.1. A *perfect Bayesian equilibrium* is a triple (π, α, ϕ) satisfying the following:

1. For all $p^t \in H$, π maximizes $C(p^t;\hat{\pi},\alpha_\ell,\alpha_h,\phi)$ among all seller strategies $\hat{\pi}$.
2. For all $p^t \in H$ and $\theta \in \{\ell, h\}$, α_θ maximizes $V_\theta(p^t;\pi,\hat{\alpha}_\theta)$ among all buyer strategies $\hat{\alpha}_\theta$.
3. For all $p^t \in H$ and $p_t \in [0,h]$, ϕ is updated according to Bayes' rule when possible; that is,

$$\phi(p^t,p_t) = \frac{\phi(p^t)\,(1-\alpha_h(p_t|p^t))}{\phi(p^t)\,(1-\alpha_h(p_t|p^t)) + (1-\phi(p^t))\,(1-\alpha_\ell(p_t|p^t))}$$

whenever the denominator is positive. Also, $\phi(\emptyset) = \phi_0$.

3 Infinite horizon

We now study the case where $T = +\infty$, which is also the case studied in Gul et al. (1986). As the introduction explains, although the analysis of Gul et al. covers the binary-demand setting as a special case, we believe it is worthwhile to study this setting in isolation. For binary demand, our model is more general than that of Gul et al., in that it allows the seller's and buyer's discount factors to be different. In the following sections we present some preliminary results, the equilibrium characterization, and some comparative statics results (in particular, the Coase conjecture). Finally, we characterize the optimal pricing for a seller with commitment power.

3.1 Preliminary results

The lemmas below are standard in the literature on bargaining with asymmetric information. The first establishes Diamond's paradox, which states that a dynamic monopolist never sets a price below the lowest buyer's valuation. Our formulation is different from (but equivalent to) the standard one.

Lemma 3.1. *A price strictly lower than ℓ is accepted for sure in any equilibrium and history.*

Lemma 3.1 implies that, in equilibrium, the seller never offers a price $p_t < \ell$, because such a price would be accepted for sure, and so offering, say, $(p_t + \ell)/2 > p_t$ (also accepted for sure) would be a profitable deviation.

Lemma 3.2 establishes the "skimming property," which states that if a given buyer is willing to accept an (on- or off-path) offer, then a higher-valuation buyer is strictly willing to accept it. For binary demand, this follows trivially from Diamond's paradox. An implication is that, along any on- or off-path history, the seller's posterior decreases unless trade occurs for sure in a given period; that is, for any history p^t such that $p_t > \ell$, we have $\phi(p^t) \le \phi(p^{t-1})$.

Lemma 3.2. *In any equilibrium, if an on- or off-path price offer is accepted with positive probability by the ℓ-buyer, then it is accepted for sure by the h-buyer.*

3.2 Equilibrium characterization

The following theorem provides a characterization of the equilibrium behavior; it is analogous to Theorem 1 in Gul et al. (1986).

Theorem 3.1. *An equilibrium exists. There is a strictly increasing sequence $(\hat{\phi}_k, \hat{p}_k)_{k=0}^{\infty}$, with $\hat{\phi}_0 = 0$ and $\hat{p}_0 = \ell$, such that, in any equilibrium, the following hold:*

1. *If k is such that $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$, then the on-path history is $(\hat{p}_k, \hat{p}_{k-1}, ..., \hat{p}_0)$ and the corresponding beliefs are $(\phi_0, \hat{\phi}_{k-1}, ..., \hat{\phi}_0)$.*
2. *If k is such that $\phi_0 = \hat{\phi}_k$, then the seller mixes between the on-path history $(\hat{p}_k, \hat{p}_{k-1}, ..., \hat{p}_0)$, with corresponding beliefs $(\phi_0, \hat{\phi}_{k-1}, ..., \hat{\phi}_0)$, and the on-path history $(\hat{p}_{k-1}, \hat{p}_{k-2}, ..., \hat{p}_0)$, with corresponding beliefs $(\phi_0, \hat{\phi}_{k-2}, ..., \hat{\phi}_0)$.*

Theorem 3.1 characterizes equilibrium behavior in terms of a sequence $(\hat{\phi}_k, \hat{p}_k)_{k=0}^{\infty}$. This sequence can easily be computed using a system of difference equations, presented in Section 3.3. Note that if $\phi_0 \notin \{\hat{\phi}_k\}_{k=1}^{\infty}$, then the equilibrium outcome is unique, while if $\phi_0 \in \{\hat{\phi}_k\}_{k=1}^{\infty}$, then all randomizations between the two price paths described are equilibrium outcomes.

Sketch of the proof

We now sketch the proof of Theorem 3.1. Readers interested in the Coase conjecture may jump directly to Section 3.3.

The proof is by induction. Let $\hat{\phi}_1 \in [0,1]$ be the largest prior such that whenever $\phi_0 \in [0,\hat{\phi}_1)$, the seller offers ℓ in the first period with probability one in all equilibria. Note that if $\phi_0 < \hat{\phi}_1$, then the h-buyer accepts for sure any price strictly lower than $\hat{p}_1$, the price at which he is indifferent between accepting in the current period and waiting to accept ℓ in the next period (i.e., $\hat{p}_1 > \ell$ satisfies $h - \hat{p}_1 = \delta_b (h - \ell)$). Hence, if $\phi_0 \in [0,\hat{\phi}_1)$, the seller can obtain a payoff arbitrarily close to

$$\phi_0 \, \hat{p}_1 + \delta_s (1 - \phi_0) \, \ell \tag{1}$$

by offering a price slightly below $\hat{p}_1$. Since (1) must be smaller than ℓ for $\phi_0 < \hat{\phi}_1$ (because, by the definition of $\hat{\phi}_1$, there is an equilibrium where ℓ is offered in the first period), we have that $\hat{\phi}_1 \leq \hat{\phi}'_1$, where $\hat{\phi}'_1 := \frac{\ell - \delta_s \ell}{\hat{p}_1 - \delta_s \ell}$ is the prior that makes (1) equal to ℓ.

The proof then argues that $\hat{\phi}_1 \geq \hat{\phi}'_1$. To see this, we let $C^\varepsilon \geq \ell$ denote the supremum of the set of payoffs the seller obtains in equilibria featuring a first-period price strictly higher than ℓ, where the supremum is taken among all priors in $[\hat{\phi}_1, \hat{\phi}_1 + \varepsilon]$. Take a sequence $\phi_0^n \to \hat{\phi}_1$ and a corresponding sequence of equilibria with first-period prices $p_0^n > \ell$ for all n and seller payoffs converging to $C^0 := \lim_{\varepsilon \searrow 0} C^\varepsilon$, which exist by the definitions of $\hat{\phi}_1$ and C^ε. If there is an increasing sequence $m^n \to \infty$ with $\phi_1^{m^n}(p_0) \geq \hat{\phi}_1$ for all n, then the probability of transaction in period 0 vanishes as $n \to 0$; this implies $C^0 \leq \delta_s C^0$, contradicting that $C^0 \geq \ell$. Hence, it must be that the on-path second-period price is ℓ if n is large enough, which implies that the first-period price offered by the seller is no larger than $\hat{p}_1$. It then follows that if ϕ_0 is higher than but close to $\hat{\phi}_1$, expression (1) must be higher than ℓ, and so $\hat{\phi}_1 \geq \hat{\phi}'_1$.

We have deduced that expression (1) is equal to ℓ when $\phi_0 = \hat{\phi}_1$, that is,

$$\hat{\phi}_1 = \left(1 + \tfrac{1-\delta_b}{1-\delta_s} \, \tfrac{h-\ell}{\ell}\right)^{-1} . \tag{2}$$

The fact that, whenever $\phi_0 < \hat{\phi}_1$, the seller sets a price equal to ℓ in all equilibria serves as an anchor for the inductive construction of the equilibrium behavior. In particular, this fact implies that if the seller's posterior is strictly below $\hat{\phi}_1$ at some (on- or off-path) history in some equilibrium (i.e., not necessarily in the first period), then the seller offers ℓ for sure at this history, while if the seller's posterior is strictly above $\hat{\phi}_1$, she offers a price no lower than $\hat{p}_1$.

We now carry out the second step of the inductive argument. (This and all subsequent steps are similar to the first.) Let $\hat{\phi}_2 \geq \hat{\phi}_1$ be the maximal posterior such that, whenever $\phi_0 \in (\hat{\phi}_1, \hat{\phi}_2)$, the seller offers $\hat{p}_1$ with probability one in the first period in all equilibria. Let $\hat{p}_2$ be the price at which the h-buyer is indifferent between accepting now and waiting to accept $\hat{p}_1$ in the next period (i.e., $\hat{p}_2$ satisfies $h - \hat{p}_2 = \delta_b (h - \hat{p}_1)$). Now, if $\phi_0 \in (\hat{\phi}_1, \hat{\phi}_2)$ and the seller offers a price p_0 slightly lower than $\hat{p}_2$, the probability that the buyer will accept p_0 must be such that the seller's posterior in the second period is $\hat{\phi}_1$. Indeed, if the second-period posterior is strictly above $\hat{\phi}_1$, then the h-buyer strictly benefits from accepting p_0, which leads to a second-period posterior strictly below $\hat{\phi}_1$, a contradiction. Similarly, if the second-period posterior is strictly below $\hat{\phi}_1$, then the h-buyer strictly benefits from rejecting p_0, which leads to a second-period posterior strictly above $\hat{\phi}_1$, again a contradiction. Therefore, if $\phi_0 \in (\hat{\phi}_1, \hat{\phi}_2)$, the seller can obtain a payoff arbitrarily close to[3]

$$\tfrac{\phi_0 - \hat{\phi}_1}{1 - \hat{\phi}_1}\, \hat{p}_2 + \tfrac{1-\phi_0}{1-\hat{\phi}_1}\, \delta_s \left(\hat{\phi}_1 \hat{p}_1 + \delta_s (1 - \hat{\phi}_1) \ell \right) . \tag{3}$$

This is no larger than the payoff the seller obtains from offering a price slightly below $\hat{p}_1$ (namely (1)) only if $\phi_0 \leq \hat{\phi}_2'$, where $\hat{\phi}_2'$ is the prior that makes (1) and (3) equal; hence $\hat{\phi}_2 \leq \hat{\phi}_2'$. On the other hand, we prove that $\hat{\phi}_2 \geq \hat{\phi}_2'$ by showing that, when ϕ_0 is higher than but close to $\hat{\phi}_2$, the seller offers $\hat{p}_2$ for sure in the first period in all equilibria; therefore, $\hat{\phi}_2 = \hat{\phi}_2'$.

Remark 3.1 (Comparison with the proof in Gul et al. (1986)). Some of the steps in our proof of Theorem 3.1 resemble those in the proof of Theorem 1 in Gul et al. (1986). (For example, Gul et al. also begin by proving Diamond's paradox, and they also show if the residual demand is small enough, then the lowest valuation is offered in the first period in all equilibria.) However, we have not been able to find a clear mapping between the two proofs. The proof in Gul et al. (1986) is considerably more involved than ours, requiring a significant amount of notation and several new concepts and intermediate results to address the case of general demand.

3.3 Comparative statics

In this section, we provide some comparative statics results. The first is the classical Coase conjecture: As $\Delta \to 0$, the physical time it takes for trade to happen shrinks to 0, and the transaction price tends to ℓ. The second result is that, as the seller becomes more patient, Coasian forces vanish, and she extracts all trade surplus from the buyer. The third result is that

[3] Note that if the prior is ϕ_0 and the posterior in the second period is equal to $\hat{\phi}_1$, then Bayes' rule specifies that $\hat{\phi}_1 = \frac{\phi_0 \alpha_h(\hat{p}_1)}{\phi_0 \alpha_h(\hat{p}_1) + 1 - \phi_0}$, which implies that the probability of trade in the first period is $\phi_0 \alpha_h(\hat{p}_1) = \frac{\phi_0 - \hat{\phi}_1}{1 - \hat{\phi}_1}$.

Coasian forces also vanish as the seller becomes more convinced that the buyer's valuation is high, and in this case the equilibrium outcome converges to the perfect information outcome.

The frequent-offers limit

We first consider the limit as the length of each period vanishes, that is, as $\Delta \to 0$. We will prove that the Coase conjecture holds, independently of the values of r_s and r_b.

To begin, we provide an explicit construction of the sequence $(\hat{\phi}_k, \hat{p}_k)_{k=0}^{\infty}$ described in Theorem 3.1. Our construction uses an auxiliary sequence $(\hat{C}_k)_{k=0}^{\infty}$, where $\hat{C}_k$ represents the seller's equilibrium payoff when the posterior is $\hat{\phi}_k$. Set $(\hat{\phi}_0, \hat{p}_0, \hat{C}_0) := (0, \ell, \ell)$. For each $k \geq 1$, we have

$$\hat{p}_k = (1-\delta_b)\, h + \delta_b\, \hat{p}_{k-1} \text{ , and} \tag{4}$$

$$\hat{C}_k = \frac{\hat{\phi}_k - \hat{\phi}_{k-1}}{1-\hat{\phi}_{k-1}}\, \hat{p}_k + \frac{1-\hat{\phi}_k}{1-\hat{\phi}_{k-1}}\, \delta_s\, \hat{C}_{k-1} = \frac{\hat{\phi}_k - \hat{\phi}_{k-1}}{1-\hat{\phi}_{k-1}}\, \hat{p}_{k-1} + \frac{1-\hat{\phi}_k}{1-\hat{\phi}_{k-1}}\, \hat{C}_{k-1} \text{ .} \tag{5}$$

Given $(\hat{\phi}_{k-1}, \hat{p}_{k-1}, \hat{C}_{k-1})$, equations (4) and (5) uniquely define $(\hat{\phi}_k, \hat{p}_k, \hat{C}_k)$. Equation (4) expresses the requirement that the h-buyer is indifferent between accepting $\hat{p}_k$ in the given period and accepting $\hat{p}_{k-1}$ the next period (recall the iterative construction of $\hat{p}_k$ described above). Equation (5) expresses the fact that, when the prior is $\hat{\phi}_k$, the seller is indifferent in equilibrium between offering $\hat{p}_k$ and offering $\hat{p}_{k-1}$ (recall the statement of Theorem 3.1 and the sketch of the proof). Indeed, the first equality in (5) says that the seller obtains her equilibrium payoff of $\hat{C}_k$ when the posterior is $\hat{\phi}_k$ if she offers $\hat{p}_k$, and the probability that the h-buyer will accept $\hat{p}_k$ is such that the next period's posterior and continuation value are $\hat{\phi}_{k-1}$ and $\hat{C}_{k-1}$, respectively. The second equality in (5) says that the seller can obtain the same payoff by charging $\hat{p}_{k-1}$, in which case the next period's posterior is $\hat{\phi}_{k-2}$.[4]

The following result, which is analogous to Theorem 3 in Gul et al. (1986), shows that the Coase conjecture holds in the binary-demand setting.

Corollary 3.1 (Coase conjecture). *For any sequence $\Delta^n \to 0$ and corresponding sequence of equilibria, the expected price offered in the initial period converges to ℓ, and the expected physical time it takes for the price to reach ℓ shrinks to 0.*

---

[4] By charging $\hat{p}_{k-1}$, the seller obtains $\frac{\hat{\phi}_k - \hat{\phi}_{k-2}}{1-\hat{\phi}_{k-2}}\, \hat{p}_{k-1} + \frac{1-\hat{\phi}_k}{1-\hat{\phi}_{k-2}}\, \delta_s\, \hat{C}_{k-2}$. Using the expression for $\hat{C}_{k-1}$, it is easy to see that this is equal to the expression on the right-hand side of the second equality in equation (5).

The proof of Corollary 3.1 is simple. Note that from equations (4) and (5) we have

$$1-\hat{\phi}_k = \frac{h-\hat{C}_k}{h-\frac{\delta_s-\delta_b}{1-\delta_b}\hat{C}_{k-1}}\,(1-\hat{\phi}_{k-1}) = \prod_{k'=1}^{k}\frac{h-\hat{C}_{k'}}{h-\frac{\delta_s-\delta_b}{1-\delta_b}\hat{C}_{k'-1}}\ ,$$

where the second equality follows from iterated use of the first equality. Since the cost to the h-buyer of waiting k periods for the price to reach ℓ vanishes as $\Delta\to 0$, it is clear from equation (4) that $\hat{p}_k\to\ell$ as $\Delta\to 0$ for each fixed k. This implies that each $\hat{C}_k$ tends to ℓ as $\Delta\to 0$. We can then use the previous expression to obtain the limit of the belief threshold for each k as $\Delta\to 0$:

$$\lim_{\Delta\to 0}\hat{\phi}_k = 1-\left(1+\tfrac{r_s}{r_b}\tfrac{\ell}{h-\ell}\right)^{-k}\ . \tag{6}$$

Equation (6) illustrates the severity of the seller's commitment problem: Even when Δ is small, she sells with a significantly high probability to the h-buyer just before selling to the ℓ-buyer. In other words, for each ϕ_0, the number of periods that elapse before she sells to the ℓ-buyer is uniformly bounded across all $\Delta>0$; hence, the physical time it takes for trade to occur in equilibrium vanishes as $\Delta\to 0$. Since $\hat{p}_k\to\ell$ as $\Delta\to 0$ for all k, this implies that the seller sells to the h-buyer at a price very close to ℓ, and so her payoff tends to ℓ.[5]

For further intuition, let t denote the period in which trade occurs in equilibrium. Then $\phi_{t-1}\in[\hat{\phi}_1,\hat{\phi}_2]$; that is, in period $t-1$, the seller is willing to trade with the h-buyer at price $\hat{p}_1$ and to postpone trading with the ℓ-buyer for one period. If the seller deviates by offering ℓ in period $t-1$, she incurs both a benefit (from selling to the ℓ-buyer one period earlier) and a loss (from selling to the h-buyer at price ℓ instead of at $\hat{p}_1=\ell+(1-\delta_b)(h-\ell)>\ell$). The payoff difference from this deviation is therefore

$$\overbrace{(1-\phi_{t-1})(1-\delta_s)\,\ell}^{\text{benefit}} - \overbrace{\phi_{t-1}(1-\delta_b)(h-\ell)}^{\text{loss}}\ .$$

Because the discount factors of the seller and the buyer converge to 1 at the same rate as $\Delta\to 0$ (note that $\lim_{\Delta\to 0}\frac{1-\delta_s}{1-\delta_b}=\frac{r_s}{r_b}$), ϕ_{t-1} cannot be close to 0 for the previous equation to be non-positive when Δ is small. Hence, price discrimination is beneficial only if ϕ_{t-1} is significantly far from 0 even when Δ is small; formally, we have $\lim_{\Delta\to 0}\hat{\phi}_1>0$. As we have argued, this implies that there must be a large equilibrium probability of transaction with the h-buyer in

[5]Formally, for any $\phi_0\in(0,1)$, there is some $k^*(\phi_0)\in\mathbb{N}$ (independent of Δ) such that $\phi_0<\lim_{\Delta\to 0}\hat{\phi}_{k^*(\phi_0)}$. Then, as $\Delta\to 0$, the seller's equilibrium payoff is no larger than $\lim_{\Delta\to 0}\hat{C}_{k^*(\phi_0)}=\ell$.

the periods before the game ends (at prices close to ℓ), which leads to the Coase conjecture.

Patient seller

An important feature of a Coasian setting like ours is that the seller is in competition with her future selves, even when she is more patient than the buyer. This is reinforced if the buyer believes so: If the buyer believes that tomorrow's price will be low, then the seller may be induced to offer a low price today. The following result establishes that when the seller is sufficiently patient, such self-fulfilling prophecies do not occur in equilibrium, so Coasian dynamics are averted and the seller obtains the full trade surplus. (Note that the limit as the buyer becomes more patient is trivial: Trade occurs immediately at price ℓ.)

Corollary 3.2. *For any sequence $r_s^n \to 0$ and corresponding sequence of equilibria, the seller's equilibrium payoff converges to $\phi_0 h + (1-\phi_0)\ell$.*

To see why Corollary 3.2 holds, note that if $\phi_0 > \hat{\phi}_k$, the seller's payoff from first offering a price slightly below $\hat{p}_k$ and then offering ℓ in the next period is approximately equal to

$$\tfrac{\phi_0 - \hat{\phi}_{k-1}}{1-\hat{\phi}_{k-1}}\, \hat{p}_k + \tfrac{1-\phi_0}{1-\hat{\phi}_{k-1}}\, \delta_s\, \ell \ . \tag{7}$$

This expression is close to the total trade surplus (which is an upper bound on the seller's equilibrium payoff) when δ_s is close to 1, $\hat{\phi}_{k-1}$ is close to 0, and $\hat{p}_k$ is close to h. Since $\hat{p}_k$ does not depend on r_s, it is readily seen from equation (5) that, for each fixed k, $\hat{\phi}_k - \hat{\phi}_{k-1} \to 0$ as $r_s \to 0$. Intuitively, Coasian forces weaken when r_s is small; hence, a patient seller finely screens the buyer before offering price ℓ. For the same reason, $\hat{\phi}_1 \to 0$ as $r_s \to 0$ (see equation (2)). As a result, $\lim_{r_s \to 0} \hat{\phi}_k = 0$ for all k: For each ϕ_0, the value of k such that $\phi_0 \in [\hat{\phi}_k, \hat{\phi}_{k+1})$ increases towards infinity as r_s shrinks. Since the buyer's discount factor is fixed, $\hat{p}_k \to h$ as $k \to \infty$ (from equation (4)). It is then clear that if r_s is small enough, the seller's equilibrium payoff must be close to the total trade surplus.[6,7]

[6] Indeed, for every $\varepsilon > 0$ there is some $\bar{r}_s > 0$ such that if $r_s < \bar{r}_s$, then there is some $k(r_s)$ satisfying $\hat{\phi}_{k(r_s)-1} < \varepsilon$ and $\hat{p}_{k(r_s)-1} > h - \varepsilon$. Then, for any ϕ_0, if r_s is close enough to 0, the seller can obtain a payoff close to the trade surplus by offering $\hat{p}_{k(r_s)-1}$ in the first period and ℓ in the second period.

[7] Note that the Coasian logic of competition with one's future selves applies before trade occurs: For example, the seller offers price $\hat{p}_1$ (which is smaller than h and independent of r_s) in the period before she offers ℓ. However, since $\lim_{r_s \to 0} \hat{\phi}_1 = 0$, when r_s is small, the seller offers $\hat{p}_1$ only when the posterior is very close to 0; hence Coasian forces have little effect on her equilibrium payoff.

A similar result holds in the limit as $\ell \to 0$, that is, in the limit where the "gap" between the seller's cost and the lowest buyer's valuation vanishes. That is, for each $\phi_0 \in (0,1)$ and $\varepsilon > 0$, there is some $\overline{\ell}_\varepsilon > 0$ such that if $\ell < \overline{\ell}_\varepsilon$, then the equilibrium payoff of the seller (in all equilibria) is higher than $\phi_0 h - \varepsilon$. This result is somehow consistent with Ausubel and Deneckere (1989)'s "folk theorem" for the "no gap" case: In our setting, for any $C \in [0, \phi_0 h]$, there are two sequences $r_s^n \to 0$ and $\ell^n \to 0$ such that the seller's equilibrium payoff in a corresponding sequence of equilibria converges to C.

The perfect information limit

The next result concerns the limit as $\phi_0 \to 1$, that is, the limit in which the buyer's valuation is known to be high (note that the limit as $\phi_0 \to 0$ is trivial). This limit can be interpreted as a perturbed version of a perfect-information bargaining model in which the buyer is known to have valuation h, but there is a small-probability behavioral "tough buyer" who only accepts offers weakly lower than ℓ (in the spirit of reputational bargaining; see Abreu and Gul, 2000). The following result states that in our setting, as the likelihood that the buyer is tough vanishes, the seller extracts the full trade surplus from the buyer. In other words, the equilibrium outcome of the asymmetric-information game converges to the outcome of the perfect-information game (where trade occurs immediately at price h) as the private information vanishes.

Corollary 3.3. *For any sequence $\phi_0^n \to 1$ and corresponding sequence of equilibria, the seller's equilibrium payoff converges to h.*

The intuition for Corollary 3.3 is as follows. First, because the buyer discounts the future, $\hat{p}_k \to h$ as $k \to \infty$. Second, for every k and $\phi_0 > \hat{\phi}_k$, the seller's equilibrium payoff is no smaller than the expression (7), which converges to $\hat{p}_k$ as $\phi_0 \to 1$. Therefore, the seller's payoff converges to h as $\phi_0 \to 1$.

3.4 Seller commitment

In this section, we briefly present the optimal strategy of a seller with commitment power. We do so for completeness and to better understand the severity of the commitment problem in our previous analysis.

To level the field, we focus on commitment to a pricing strategy instead of a general mechanism. We consider a game in which first the seller commits to a pricing strategy, and then the buyer chooses the best response. We let $\phi^* := \ell/h$.

Theorem 3.2. *When the seller has commitment power, the equilibrium outcomes are as follows:*

1. *If $\delta_s \leq \delta_b$ then there is trade only in period 0. Trade occurs at price ℓ if $\phi_0 < \phi^*$, at price h if $\phi_0 > \phi^*$, or at a price drawn from $\{\ell, h\}$ if $\phi_0 = \phi^*$.*
2. *If $\delta_s > \delta_b$ then the outcome has the following form:*
   (a) If $\phi_0 < \hat{\phi}_1$, then trade occurs immediately at price ℓ.
   (b) If $\phi_0 > \hat{\phi}_1$, then the seller offers some price $p_0 \in (\ell, h)$ in period 0, which the h-buyer accepts; moreover, there is some $\bar{t} > 0$ such that the ℓ-buyer trades for sure in period $\bar{t}$ or period $\bar{t}+1$, at price ℓ.
   (c) When $\phi_0 = \hat{\phi}_1$, the seller mixes between (i) offering ℓ in the first period and (ii) offering $\hat{p}_1$ in the first period and ℓ in the second period.

The first part of Theorem 3.2 is as expected: If the seller is at least as impatient as the buyer, she commits not to intertemporally price-discriminate and obtains the static monopolistic payoff. When $\delta_s = \delta_b$, this follows from the classical result in Stokey (1979). To see why it also holds when $\delta_s < \delta_b$, assume for the sake of contradiction that, for some δ_s and δ_b with $\delta_s < \delta_b$, the seller has a (weakly) optimal strategy involving price discrimination. Then a seller with discount factor equal to δ_b could use the same strategy to obtain a higher payoff than she would get from the optimal non-price-discriminating strategy – which contradicts the result of Stokey (1979). Therefore, while commitment power does not enable an impatient seller to capture all trade surplus, it benefits her when $\phi_0 > \phi^*$: In this case, whenever $\delta_s \leq \delta_b$, the seller would prefer to commit not to price-discriminate, so she is strictly worse off if she lacks commitment power.

The second part of Theorem 3.2 establishes that when the seller is more patient than the buyer, commitment power enables her to price-discriminate, taking advantage of the buyer's high delay cost. (For analogous results, see Fudenberg and Tirole, 1983, and Landsberger and Meilijson, 1985.) Unlike when $\delta_s \leq \delta_b$, the set of priors for which the seller offers ℓ is now independent of her commitment power (it equals $[0, \hat{\phi}_1]$). For higher priors, a seller with commitment power sells to the h-buyer earlier, at a high price, and sells to the ℓ-buyer significantly later; a seller without commitment power sells to the h-buyer later, at a low price, and sells to the ℓ-buyer without significant further delay (at least when Δ is small). From the proof of Theorem 3.2, it is easy to see that, generically in ϕ_0, the time at which the seller with commitment power offers ℓ is deterministic (i.e., the seller offers ℓ at some time $\bar{t}$ for sure).

4 Finite horizon

We now consider the finite-horizon case: We assume the game ends in period $T \in \mathbb{Z}_+$. This setting is a binary version of the model in Fuchs and Skrzypacz (2013), which assumes that the buyer's valuation is distributed according to a power distribution; hence the distribution is absolutely continuous and exhibits no gap. The setting is also a generalization of the two-period binary-demand setting of Fudenberg and Tirole (1983) to an arbitrarily long horizon.

4.1 Equilibrium characterization

The following result is analogous to Theorem 3.1 in that it provides a full characterization of equilibrium behavior.

Theorem 4.1. *An equilibrium exists. There exists an increasing sequence $(k_T, \overline{\phi}_T)_{T=0}^{\infty}$, with $\overline{\phi}_T \geq \hat{\phi}_{k_T}$ for all T and $(k_0, \overline{\phi}_0) = (0, \phi^*)$, such that, in any equilibrium, the following hold:*

1. *If $\phi_0 > \overline{\phi}_T$, then the on-path history is $(h, h, ..., h)$ and the corresponding beliefs are $(\phi_0, \overline{\phi}_{T-1}, ..., \overline{\phi}_0)$.*
2. *If $\phi_0 < \overline{\phi}_T$, then we have the following:*
   (a) If $\phi_0 \in (\hat{\phi}_{k_T}, \overline{\phi}_T)$, then the on-path history is $(\hat{p}_{k_T}, \hat{p}_{k_T - 1}, ..., \ell)$ and the corresponding beliefs are $(\phi_0, \hat{\phi}_{k_T - 1}, ..., \hat{\phi}_0)$.
   (b) Otherwise, the equilibrium is as specified in Theorem 3.1.
3. *If $\phi_0 = \overline{\phi}_T$, then the seller randomizes between the paths in parts 1 and 2.*[8]

Theorem 4.1 establishes that when there is a deadline in bargaining, two types of equilibrium outcome are possible. The first type includes "Coasian" outcomes like those described in Theorem 3.1: The equilibrium prices follow a decreasing sequence $(\hat{p}_k, \hat{p}_{k-1}, ..., \ell)$, and in particular the seller may offer ℓ before the deadline is reached. The second type is the "high-price" outcome, in which the seller maintains a "tough" position by offering h at all times. In this case, because the h-buyer does not obtain any surplus from trading, he is indifferent between

[8] That is, if $\phi_0 = \overline{\phi}_T > \hat{\phi}_{k_T}$, then she randomizes between the paths in parts 1 and 2(a), and if $\phi_0 = \overline{\phi}_T = \hat{\phi}_{k_T}$, then she randomizes between the paths in parts 1 and 2(b).

accepting and rejecting the price h in any period, and he accepts it with positive probability in all periods before the deadline.

It is not difficult to see that $\lim_{T\to\infty}\overline{\phi}_T = 1$ and $\lim_{T\to\infty}\overline{k}_T = \infty$, so the equilibrium outcome converges to the outcome of the infinite-horizon case as the horizon increases. That is, for a given ϕ_0, the equilibrium outcome coincides with that in Theorem 3.1 if T is large enough.

Proof sketch

The proof of Theorem 4.1 differs from that of Theorem 3.1 in several respects. The finite horizon means we can analyze the equilibria using backward induction from the last period, which simplifies the initial step of the iterative argument. However, the nonstationarity of the finite-horizon setting complicates the analysis. An additional complication is that, as described above (and unlike in Fuchs and Skrzypacz, 2013), different types of equilibria may arise depending on the prior and the time horizon – some exhibiting Coasian dynamics, with trade occurring for sure before the deadline, and some in which the seller is tough, offering only the high price.

The set of equilibria is constructed as follows. For the first step in our inductive argument, we study a model with $T = 0$, that is, a one-period game. In all equilibria of this game, the seller offers ℓ if $\phi_0 < \phi^*$ and h if $\phi_0 > \phi^*$; if $\phi_0 = \phi^*$, she mixes between these two offers (recall that $\phi^* \equiv \ell/h$). Hence, $\overline{\phi}_0 = \phi^*$ and $k_0 = 0$.

For the second step, consider a model with $T = 1$. We first observe that if $\phi_0 > \phi^*$, the seller can ensure a payoff arbitrarily close to

$$\tfrac{\phi_0-\phi^*}{1-\phi^*}\, h + \delta_s \tfrac{1-\phi_0}{1-\phi^*}\, \phi^*\, h \tag{8}$$

by charging a price slightly below h in periods 0 and 1. Indeed, suppose the seller (on or off path) offers a price p_0 slightly below h in period 0. In that case, the posterior in period 1, must be equal to ϕ^*. The argument is as follows. If, in equilibrium, the second-period posterior is strictly above ϕ^* (so the h-buyer rejects p_0 with positive probability in the first period), the second-period price is h; but this gives the h-buyer a strict incentive to accept $p_0 < h$ in the first period, a contradiction. If instead the second-period posterior is strictly below ϕ^* (so the h-buyer accepts p_0 with positive probability in the first period), the price in the second period is ℓ; but this gives the h-buyer a strict incentive to reject p_0 in the first period if p_0 is close enough to h, again a contradiction. Therefore the buyer must be indifferent between accepting and rejecting p_0, which means the seller mixes between ℓ and h in period 1; hence $\phi_1(p_0)=\phi^*$.

We then define $\overline{\phi}_1$ as the unique solution of

$$\tfrac{\overline{\phi}_1-\phi^*}{1-\phi^*}\,h+\delta_{\mathrm{s}}\,\tfrac{1-\overline{\phi}_1}{1-\phi^*}\,\phi^*\,h=\max\left\{\,\ell\,,\overline{\phi}_1\,\hat{p}_1+\delta_{\mathrm{s}}\,(1-\overline{\phi}_1)\,\ell\,\right\}\;;$$

that is, if $\phi_0=\overline{\phi}_1$, then the seller is indifferent between, on the one hand, offering a price arbitrarily close to h (and obtaining a payoff arbitrarily close to the left-hand side) and, on the other hand, following a Coasian equilibrium by either offering $\hat{p}_1$ (and then ℓ in period 1) or offering ℓ. If $\overline{\phi}_1<\hat{\phi}_1$, then $\bar{k}_1=0$, and in period 0 in any equilibrium, the seller either offers ℓ (if $\phi_0<\overline{\phi}_1$) or h (if $\phi_0>\overline{\phi}_1$), or randomizes between these two prices (if $\phi_0=\overline{\phi}_1$). If $\overline{\phi}_1>\hat{\phi}_1$, then $\bar{k}_1=1$, and in period 0 in any equilibrium, the seller either offers ℓ (if $\phi_0<\overline{\phi}_1$), or $\hat{p}_1$ (if $\hat{\phi}_1<\phi_0<\overline{\phi}_1$), or h (if $\phi_0>\overline{\phi}_1$), or uses a corresponding mixing (whenever $\phi_0\in\{\hat{\phi}_1,\overline{\phi}_1\}$). The non-generic case where $\overline{\phi}_1=\hat{\phi}_1$ is equivalent to the case where $\overline{\phi}_1<\hat{\phi}_1$, except that $\bar{k}_1=1$, and if $\phi_0=\overline{\phi}_1$, then the seller may mix between ℓ, $\hat{p}_1$, and h in period 0.

The characterization of equilibrium behavior for $T\geq 2$ is analogous.

4.2 The frequent-offers limit

We now consider the frequent-offers limit, as we did in Section 3.3 for the infinite-horizon case. In particular, we are interested in the limit as the length of each bargaining period gets smaller (i.e., $\Delta\to 0$), while the horizon converges to some fixed $\overline{\tau}>0$. Thus, for each $\Delta>0$, we consider the finite-horizon model with $T_\Delta:=\max\{t|t\,\Delta<\overline{\tau}\}$ periods, each of length Δ. As $\Delta\to 0$ we have $T_\Delta\to\infty$, while the physical time available until the deadline, $T_\Delta\,\Delta$, converges to $\overline{\tau}$.[9]

Before calculating the limit of $\overline{\phi}_{T_\Delta}$ as $\Delta\to 0$, which we denote by $\overline{\phi}_{\overline{\tau}}$, we make the following observation (recall that, for a given $\Delta>0$, $\overline{\phi}_{T_\Delta}\in[\phi^*,1)$ is such that, in the T_Δ-period model, the seller offers h in all periods if $\phi_0>\overline{\phi}_{T_\Delta}$, while the equilibrium is Coasian if $\phi_0<\overline{\phi}_{T_\Delta}$). Note that if $\phi_0>\phi^*$ and $\overline{\tau}$ is not too high, then we must have $\overline{\phi}_{\overline{\tau}}\in(\phi^*,1)$. Indeed, $\overline{\phi}_{\overline{\tau}}=1$ would imply that, for all $\phi_0\in(0,1)$, the seller's equilibrium payoff is close to ℓ for Δ small enough. However, by waiting until the deadline and then offering h, the seller can obtain a payoff of $e^{-r_{\mathrm{s}}\overline{\tau}}\,\phi_0\,h$, which is larger than ℓ if $\phi_0>\phi^*$ and $\overline{\tau}$ is small enough. Similarly, if $\overline{\phi}_{\overline{\tau}}=\phi^*$, then there is (almost) no trade between the second period and the deadline if Δ is small enough (recall that, if $\phi^*>\overline{\phi}_{T_\Delta}$,

[9]It is not difficult to see that if we let $\Delta\to 0$ while T remains fixed, the equilibrium transaction price approaches that of the static model (with $T=0$) in equilibrium.

then $\overline{\phi}_t \geq \phi^*$ for all t). This implies that the seller's continuation value after the first period is close to $e^{-r_s \overline{\tau}}\,\ell$, but then she can profitably deviate by offering ℓ in the second period. The following result establishes that, in fact, $\overline{\phi}_{\overline{\tau}} \in (\phi^*, 1)$ for all $\overline{\tau} \in (0, +\infty)$.

Corollary 4.1. *We have* $\lim_{\Delta\to 0} \overline{\phi}_{T_\Delta} = \overline{\phi}_{\overline{\tau}}$, *where*

$$\overline{\phi}_{\overline{\tau}} = 1 - (1-\phi^*)\, e^{-r_s \frac{\ell}{h-\ell} \overline{\tau}} \ . \tag{9}$$

Equation (9) is obtained as follows. Note first that if $\phi_0 > \overline{\phi}_{T_\Delta}$ for a small $\Delta > 0$, the continuation payoff after the first period should be close to ℓ. Indeed, the posterior in the second period is $\overline{\phi}_{T_\Delta - 1}$, and from the proof of Theorem 4.1 (and from the sketch of the proof above), we can see that the continuation payoff in the second period is equal to

$$\frac{\overline{\phi}_{T_\Delta-1} - \hat{\phi}_{\overline{k}_{T_\Delta-1}-1}}{1-\hat{\phi}_{\overline{k}_{T_\Delta-1}-1}}\, \hat{p}_{\overline{k}_{T_\Delta-1}-1} + \frac{1-\overline{\phi}_{T_\Delta-1}}{1-\hat{\phi}_{\overline{k}_{T_\Delta-1}-1}}\, \delta_s\, \hat{C}_{\overline{k}_{T_\Delta-1}-2} \ ,$$

since the seller is indifferent between offering h and offering $\hat{p}_{\overline{k}_{T_\Delta-1}-1}$, in which case the next period's posterior would be $\hat{\phi}_{\overline{k}_{T_\Delta-1}-1}$. A straightforward argument similar to the one used to prove the Coase conjecture (Corollary 3.1) shows that both

$$\lim_{\Delta\to 0} \hat{p}_{\overline{k}_{T_\Delta-1}-1} = \ell \text{ and } \lim_{\Delta\to 0} \hat{C}_{\overline{k}_{T_\Delta-1}-2} = \ell \ ,$$

because the continuation play is Coasian after the seller offers $\hat{p}_{\overline{k}_{T_\Delta-1}-1}$.[10] That is, if $\phi_0 > \overline{\phi}_{\overline{\tau}}$ and $\Delta > 0$ is small enough, the seller sells with probability approximately $\frac{\phi_0 - \overline{\phi}_{\overline{\tau}}}{1-\overline{\phi}_{\overline{\tau}}}$ in the first period at price h, and her continuation payoff is approximately ℓ. Hence, as $\Delta \to 0$, her equilibrium payoff converges to

$$\frac{\phi_0 - \overline{\phi}_{\overline{\tau}}}{1-\overline{\phi}_{\overline{\tau}}}\, h + \frac{1-\phi_0}{1-\overline{\phi}_{\overline{\tau}}}\, \ell \ . \tag{10}$$

After the initial trade burst, the seller's continuation payoff stays close to ℓ. Therefore, in equilibrium, the h-buyer accepts h at a slow rate so that the posterior at physical time τ is equal to $\overline{\phi}_{\overline{\tau}-\tau}$ (note that $\overline{\tau} - \tau$ is the physical time remaining until the deadline). This means that for

[10] Note that the claim that $\lim_{\Delta\to 0} \hat{p}_{\overline{k}_{T_\Delta-1}-1} = \ell$ and $\lim_{\Delta\to 0} \hat{C}_{\overline{k}_{T_\Delta-1}-2} = \ell$ does not follow directly from Corollary 3.1, because $\overline{k}_{T_\Delta - 1}$ could potentially increase towards infinity as $\Delta \to 0$. However, equation (6) shows that that is not the case: For a fixed ϕ_0, trade occurs for sure within a bounded number of periods (with a bound independent of Δ) under the Coase outcome.

$\varepsilon > 0$ small enough,

$$\ell = \frac{\overline{\phi}_{\overline{\tau}} - \overline{\phi}_{\overline{\tau}-\varepsilon}}{1-\overline{\phi}_{\overline{\tau}-\varepsilon}} h + e^{-r_s \varepsilon} \frac{1-\overline{\phi}_{\overline{\tau}}}{1-\overline{\phi}_{\overline{\tau}-\varepsilon}} \ell + o(\varepsilon) \;\Rightarrow\; r_s \ell = \frac{h-\ell}{1-\overline{\phi}_{\overline{\tau}}} \frac{\mathrm{d}}{\mathrm{d}\overline{\tau}} \overline{\phi}_{\overline{\tau}} \,. \tag{11}$$

Hence the probability of trade per unit of physical time is $\lim_{\varepsilon \to 0} \frac{\overline{\phi}_{\overline{\tau}} - \overline{\phi}_{\overline{\tau}-\varepsilon}}{1-\overline{\phi}_{\overline{\tau}-\varepsilon}} = r_s \ell / h$. The solution to the differential equation on the right side of (11) with boundary condition $\overline{\phi}_0 = \phi^*$ is (9). Therefore, if $\phi_0 > \overline{\phi}_{\overline{\tau}}$, there is an initial trade burst (where the posterior jumps from ϕ_0 to $\overline{\phi}_{\overline{\tau}}$), then a constant rate of trade until the deadline (as the posterior slowly declines, being equal to $\overline{\phi}_{\overline{\tau}-\tau}$ at physical time τ), and finally a trade burst at the deadline (where trade occurs with probability ϕ^*).[11]

For a fixed $\overline{\tau}$, the range of priors for which the seller offers h in equilibrium is $[1-\overline{\phi}_{\overline{\tau}}, 1]$. It is easy to see that, as one might expect, this range shrinks if the seller becomes more impatient, if the low valuation increases, or if the high valuation decreases.

We briefly compare Corollary 4.1 with the results in Fuchs and Skrzypacz (2013) and Dilmé (2023b).[12] A first important difference is that our setting exhibits trade bursts at both the outset of bargaining and the deadline, whereas Fuchs and Skrzypacz (2013) and Dilmé (2023b) find a trade burst only at the deadline. This trade burst is the source of most of the surplus the seller obtains from trade (after that, her continuation payoff is ℓ). In fact, in both of the earlier models, the seller's equilibrium payoff is the same as her payoff from waiting until the deadline and then charging the monopolistic price. In our model, the seller has a higher payoff: It is easy to see that she obtains more than $e^{-r_s \overline{\tau}} \phi_0 h$ in equilibrium, because she has a positive probability of selling to the h-buyer at price h before the deadline. Rather paradoxically, the seller's commitment problem is worse in the previous models, where there is no gap, than in ours, where there is a gap. Note also that in both our model and that of Dilmé (2023b), the seller's payoff is independent of r_b (Fuchs and Skrzypacz, 2013, consider only the case of equal discounting).

[11] Note that the rate at which trade occurs (not conditioning on the buyer's valuation) between the two bursts must be indeed constant: From equation (11), the probability of trade between τ and $\tau + \varepsilon$ is $\frac{\ell}{h} r_s \varepsilon + o(\varepsilon)$.

[12] Recall that, like Fuchs and Skrzypacz (2013), Dilmé (2023b) studies a version of the setting of Gul et al. (1986) with a deadline. Unlike the model of Fuchs and Skrzypacz (2013), that of Dilmé (2023b) is set directly in continuous time, allows for more general absolutely continuous distribution (subject to standard regularity conditions within the no-gap case), and also allows for different discount rates.

5 Conclusions

The Coase conjecture is ubiquitous in the study of dynamic monopolists and bargaining with asymmetric information. Although the conjecture is counterintuitive, Coasian forces are present in numerous settings and can completely determine equilibrium outcomes unless strong countervailing effects are present.[13] Unfortunately, proofs of the Coase conjecture are often laden with technicalities that make it difficult to comprehend their overall logic.[14]

Our main contribution in this paper is a complete characterization of the equilibrium pricing for a monopolist facing binary demand. Although our binary-demand model is a special case of the model with general demand studied in Gul et al. (1986), we have found it fruitful to study this case separately. Because it is far more tractable than the general case, we are able to provide simple arguments that shed light on the logic behind the Coase conjecture. We also derive several new results and fully characterize the equilibria in the finite-horizon case. Many of our arguments may be adaptable to other bargaining settings.

[13] Such countervailing effects may arise from discrete demand (Bagnoli et al., 1989), adverse selection (Deneckere and Liang, 2006), capacity choice (McAfee and Wiseman, 2008), arrival of traders (Fuchs and Skrzypacz, 2010), outside options (Board and Pycia, 2014), or differentiated goods (Nava and Schiraldi, 2019). Note also that Ausubel and Deneckere (1989) present a folk theorem for the no-gap case.

[14] Example 1 in Gul et al. (1986) shows that in the uniform-demand case, stationary equilibria can be analyzed in a tractable way (see Güth and Ritzberger, 1998, for a detailed analysis). However, in this case there is no gap between the lowest buyer's valuation and the seller's cost; therefore, the example fails to illustrate why the Coase conjecture holds in all equilibria (in fact, there are non-stationary, non-Coasian equilibria in the uniform-demand case). To our knowledge, no tractable example in the gap case has been studied.

A Proofs

Proof of Lemmas 3.1 and 3.2

Proof. Even though the proofs of Lemmas 3.1 and 3.2 are standard, we include them here for completeness. To prove Lemma 3.1, let $\underline{p}$ be the infimum of all prices that are optimal for the seller in some equilibrium at some history. For the sake of contradiction, assume that $\underline{p} < \ell$. Consider a history and an equilibrium where it is optimal for the seller to offer $\underline{p} + \varepsilon$, for $\varepsilon \geq 0$ small enough that both types of the buyer strictly prefer accepting $\underline{p} + \varepsilon$ this period than $\underline{p}$ next period, that is, satisfying that $\varepsilon \in [0, (1 - \delta_{\mathrm{b}})(\ell - \underline{p}))$ (which exists by the definition of $\underline{p}$ and the assumption that $\underline{p} < \ell$). Then, the seller can profitably deviate by setting a price slightly above $\underline{p} + \varepsilon$ that keeps both types of the buyer strictly willing to accept it, a contradiction. Hence, since no price strictly lower than ℓ is offered in equilibrium, if the seller deviates and offers a price strictly lower than ℓ, such a price is accepted for sure.

Lemma 3.2 follows immediately from Lemma 3.1: If an on- or off-path price offer is accepted with positive probability by the ℓ-buyer, such price offer should be weakly lower than ℓ, but then it is accepted for sure by the h-buyer because it is equal or lower than the infimum price the seller offers under any equilibrium. □

Proof of Theorem 3.1

Proof. **Notation:** We define the correspondence $\mathcal{C}^* : [0,1] \rightrightarrows [\ell, h]$ so that, for each $\phi_0 \in [0,1]$, $\mathcal{C}^*(\phi_0)$ denotes the set of seller's equilibrium payoffs for prior ϕ_0. We will show that, in fact, there is a unique seller's equilibrium payoff for each $\phi_0 \in [0,1]$. For a property "Q" of the set of equilibria, we will use $\mathcal{C}^*(\phi_0 | Q)$ to denote the equilibrium payoffs of the among the equilibria satisfying property Q when the prior is ϕ_0.

Induction argument: We will provide a proof by induction over $k = 0, 1, ...$ Our induction hypothesis in the k-th step is the following:

> **Induction hypothesis for k:** There is a strictly increasing sequence $(\hat{\phi}_{k'}, \hat{p}_{k'})_{k'=0}^{k+1}$, with $(\hat{\phi}_0, \hat{p}_0) = (0, \ell)$, such that the following is true in all equilibria:
>
> 1. For all $\phi_0 \leq \hat{\phi}_k$, let $k' < k$ be such that $\phi_0 \in (\hat{\phi}_{k'}, \hat{\phi}_{k'+1}]$. Then:

(a) If $\phi_0 \in (\hat{\phi}_{k'}, \hat{\phi}_{k'+1})$ then the on-path history is $(\hat{p}_{k'}, \hat{p}_{k'-1}, ..., \hat{p}_0)$ and the corresponding beliefs are $(\phi_0, \hat{\phi}_{k'-1}, ..., \hat{\phi}_0)$..

(b) If $\phi_0 = \hat{\phi}_{k'}$ then the seller mixes between the on-path history $(\hat{p}_{k'}, \hat{p}_{k'-1}, ..., \hat{p}_0)$, with corresponding beliefs $(\phi_0, \hat{\phi}_{k'-1}, ..., \hat{\phi}_0)$, and the on-path history $(\hat{p}_{k'-1}, \hat{p}_{k'-2}, ..., \hat{p}_0)$, with corresponding beliefs $(\phi_0, \hat{\phi}_{k'-2}, ..., \hat{\phi}_0)$.

2. For all $\phi_0 > \hat{\phi}_{k+1}$, there is no equilibrium where the seller charges a price strictly below $\hat{p}_{k+1}$ in the first period with positive probability, and the payoff of the h-buyer is weakly below $h - \hat{p}_{k+1}$ in all equilibria.

Part 1 of the induction argument: Proof for $k = 0$. We first prove that there is a pair $(\hat{\phi}_1, \hat{p}_1)$, satisfying the properties stated in the induction hypothesis. This part illustrates the second part of the proof, which will generalize the arguments to a general k.

We let $\hat{\phi}_1 \in [0,1]$ be the highest prior satisfying that, for all $\phi_0 < \hat{\phi}_1$, trade occurs for sure in the first period at price ℓ in any equilibrium. Fix an equilibrium. Note that, because the seller never offers a price below ℓ at any history of any equilibrium (by the Diamond's paradox), if the seller charges a price strictly below $\hat{p}_1 := (1-\delta_b)\, h + \delta_b\, \ell$ (on- or off-path), the h-buyer accepts such an offer for sure. Therefore, it must be that, for all $\phi_0 < \hat{\phi}_1$,

$$\ell \geq \phi_0\, \hat{p}_1 + \delta_s\, (1-\phi_0)\, \ell \;\;\Rightarrow\;\; \phi_0 \leq \hat{\phi}'_1 := \tfrac{(1-\delta_s)\,\ell}{(1-\delta_s)\,\ell + (1-\delta_b)\,(h-\ell)} \; .$$

We then have that $\hat{\phi}_1 \leq \hat{\phi}'_1$.

For each $\varepsilon > 0$, we define

$$\overline{C}^{\varepsilon}_1 := \sup\big(\overbrace{\cup_{\phi_0 \in [\hat{\phi}_1, \hat{\phi}_1+\varepsilon]} \mathcal{C}^*(\phi_0 | \Pr(p_0 > \ell) > 0)}^{(*)}\big) \; .$$

Note that, for all $\varepsilon > 0$, $(*)$ is non-empty (by the definition of $\hat{\phi}_1$). We let $\overline{C}^0_1 := \lim_{\varepsilon \searrow 0} \overline{C}^{\varepsilon}_1$, which exists because $\overline{C}^{\varepsilon}_1$ is non-decreasing in ε. We let $(\phi^n_0)^{\infty}_{n=1}$ be a sequence decreasing toward $\hat{\phi}_1$ such that there is a corresponding sequence $(\pi^n, \alpha^n, \phi^n)^{\infty}_{n=1}$ satisfying that (i) for each n, $(\pi^n, \alpha^n, \phi^n)$ is an equilibrium when the prior is ϕ^n_0 satisfying $\pi^n(p^n_0 > \ell) > 0$, and (ii) $\lim_{n\to\infty} C^n_0 = \overline{C}^0_1$ (where C^n_0 is the seller's equilibrium payoff in the n-th equilibrium). Without loss of generality for the argument, we assume that the seller does not randomize in the first period, that is, there is a sequence $(p^n_0)^{\infty}_{n=1}$ such that $p^n_0 > \ell$ and $\pi^n(p^n_0) = 1$ for all n.[15] There are

[15]The reason is that if an equilibrium where the seller offers a given p_0 with positive probability in the first period

four possibilities:

1. Assume first, for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $\pi^{n_m}(p_1^{n_m} > \ell | p_0^{n_m}) > 0$ for all m (i.e., the seller offers a price above ℓ with positive probability in the second period after offering $p_0^{n_m}$ in the first period). This implies that $\phi_1^{n_m}(p_0^{n_m}) \in [\hat{\phi}_1, \phi_0^{n_m}]$ (i.e., after the seller offers $p_0^{n_m}$ (on-path), the second period's posterior is in $[\hat{\phi}_1, \phi_0^{n_m}]$). The payoff of the seller is then

$$C_0^{n_m} = \frac{\phi_0^{n_m} - \phi_1^{n_m}(p_0^{n_m})}{1-\phi_1^{n_m}(p_0^{n_m})}\, p_0^{n_m} + \frac{1-\phi_0^{n_m}}{1-\phi_1^{n_m}(p_0^{n_m})}\, \delta_{\mathrm{s}}\, C_1^{n_m}(p_0^{n_m}) \ .$$

   By assumption, the left-hand side tends to $\overline{C}_1^0$ as $m \to \infty$. The first term on the right-hand side tends to 0 because $\phi_0^{n_m} \to \hat{\phi}_1$ as $m \to \infty$, hence $\phi_1^{n_m}(p_0^{n_m}) \to \hat{\phi}_1$ as $m \to \infty$ as well. Since $C_1^{n_m}(p_0^{n_m}) \le \overline{C}_1^{\varepsilon}$ for $\varepsilon = \phi_1^{n_m}(p_0^{n_m}) - \hat{\phi}_1$, we have that

$$\overline{C}_1^0 = \lim_{m\to\infty} C_0^{n_m} = \delta_{\mathrm{s}} \lim_{m\to\infty} C_1^{n_m}(p_0^{n_m}) \le \delta_{\mathrm{s}}\, \overline{C}_1^0 \ ,$$

   which implies that $\overline{C}_1^0 \le 0$. This is a contradiction because $\overline{C}_1^0 \ge \ell > 0$.

2. Assume now, again for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $p_0^{n_m} > \hat{p}_1$ for all m. Since, from the previous result, we have that $\pi^{n_m}(p_1^{n_m} > \ell | p_0^{n_m}) = 0$ if m is large enough, we have that the h-buyer accepts with probability zero the first price, and so $\phi_1^{n_m}(p_0^{n_m}) = \phi_0^{n_m}$ for m large enough. Now, we have $C_0^{n_m} = \delta_{\mathrm{s}}\, C_1^{n_m}(p_0^{n_m})$, which again implies that $\overline{C}_1^0 \le \delta_{\mathrm{s}}\, \overline{C}_1^0$, a contradiction.

3. Assume, again for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $p_0^{n_m} \in (\ell, \hat{p}_1)$ for all m. We argued that in this case, by the Diamond's paradox, the h-buyer accepts the first offer for sure for all m. Nevertheless, the seller can profitably deviate by offering a price in $(p_0^{n_m}, \hat{p}_1)$ (which is accepted for sure by h-buyer by the Diamond's paradox), a contradiction.

4. The only possibility left is that $p_0^n = \hat{p}_1$ and $\pi^n(p_1^n > \ell | p_0^n) = 0$ for n large enough. The payoff of the seller is then

$$C_0^n = \phi_0^n\, \hat{p}_1 + (1-\phi_0^n)\, \delta_{\mathrm{s}}\, \ell \ \Rightarrow \ \overline{C}_1^0 = \hat{\phi}_1\, \hat{p}_1 + \delta_{\mathrm{s}}\, (1-\hat{\phi}_1)\, \ell \ .$$

exists, then there is an equilibrium where the seller offers p_0 for sure in the first period.

Since it must be that $\overline{C}_1^0 \geq \ell$, we have that $\hat{\phi}_1 \geq \hat{\phi}_1'$ (recall that $\hat{\phi}_1' \hat{p}_1 + \delta_s (1-\hat{\phi}_1') \ell = \ell$).

We then conclude that, since $\hat{\phi}_1 \leq \hat{\phi}_1'$ and $\hat{\phi}_1 \geq \hat{\phi}_1'$, we have $\hat{\phi}_1 = \hat{\phi}_1'$. Note finally that if $\phi_0 > \hat{\phi}_1$, then the seller has the option of offering a price slightly below $\hat{p}_1$, which ensures that the h-buyer accepts the price for sure. The arguments above show that if $\phi_0 > \hat{\phi}_1$ then it is strictly suboptimal for the seller to offer a price strictly lower than $\hat{p}_1$. Furthermore, it easily follows from the previous arguments that if $\phi_0 = \hat{\phi}_1$, all equilibria begin with a (possibly degenerated) randomization between offering ℓ and $\hat{p}_1$, and for all mixing probabilities, there is an equilibrium with such a mixing probability in the first period.

Part 2 of the induction argument: Proof for $k \geq 1$. We now assume that the induction hypothesis holds for $k-1$, and we will prove it holds for k. We let $\hat{\phi}_{k+1}$ be the posterior such that, for all $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$, the seller offers $\hat{p}_k$ in the first period for sure and the continuation value for the h-buyer is $h - \hat{p}_k$ in all equilibria (we let $\hat{\phi}_{k+1} := \hat{\phi}_k$ if such a posterior does not exist). We divide this part of the proof into four subparts.

Part 2.1. Assume $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$ and fix an equilibrium (π, α, ϕ) (hence $\pi(p_0 = \hat{p}_k) = 1$). We show that the seller's equilibrium payoff is at most

$$\tfrac{\phi_0 - \hat{\phi}_{k-1}}{1-\hat{\phi}_{k-1}} \hat{p}_k + \tfrac{1-\phi_0}{1-\hat{\phi}_{k-1}} \delta_s \hat{C}_{k-1} \,, \tag{12}$$

where $\hat{C}_{k-1}$ is the seller's payoff in any equilibrium when the prior is $\hat{\phi}_{k-1}$ (which satisfies (5) and $\hat{C}_0 = \ell$). We divide the argument into three cases:

1. We first assume, for the sake of contradiction, that $\phi_1(\hat{p}_k) < \hat{\phi}_{k-1}$. In this case, by the induction hypothesis, the h-buyer's continuation payoff in the second period is at least $h - \hat{p}_{k-2}$, but then it has a strict incentive to reject $\hat{p}_k$, contradicting that $\phi_1(\hat{p}_k) < \hat{\phi}_{k-1} < \phi_0$.
2. Assume now, again for the sake of contradiction, that $\phi_1(\hat{p}_k) \in (\hat{\phi}_k, \hat{\phi}_{k+1})$. In this case, the h-buyer's continuation payoff in the second period is $h - \hat{p}_k$, but then he has a strict incentive to accept $\hat{p}_k$ in the first period, a contradiction.
3. Assume finally that $\phi_1(\hat{p}_k) \in [\hat{\phi}_{k-1}, \hat{\phi}_k]$. In this case, because the h-buyer accepts $\hat{p}_k$ with a non-degenerated probability, his continuation payoff in the second period should be $h - \hat{p}_{k-1}$, which using the induction hypothesis implies that $\pi(p_1 = \hat{p}_{k-1} | \hat{p}_k) = 1$, and hence

$\phi_2(\hat{p}_k, \hat{p}_{k-1}) = \hat{\phi}_{k-2}$. The seller's equilibrium payoff is then given by

$$\tfrac{\phi_0 - \phi_1(\hat{p}_k)}{1-\phi_1(\hat{p}_k)}\, \hat{p}_k + \tfrac{1-\phi_0}{1-\phi_1(\hat{p}_k)}\, \delta_{\text{s}} \Big(\underbrace{\tfrac{\phi_1(\hat{p}_k) - \hat{\phi}_{k-2}}{1-\hat{\phi}_{k-2}}\, \hat{p}_{k-1} + \tfrac{1-\phi_1(\hat{p}_k)}{1-\hat{\phi}_{k-2}}\, \delta_{\text{s}}\, \hat{C}_{k-2}}_{(*)} \Big) \, . \tag{13}$$

The derivative of the previous expression with respect to $\phi_1(\hat{p}_k)$ is

$$-\tfrac{(1-\phi_0)}{(1-\phi_1(\hat{p}_k))^2}\, (\hat{p}_k - \delta_{\text{s}}\, \hat{p}_{k-1}) < 0 \, .$$

Hence, the seller's equilibrium payoff is bounded by expression (13) evaluated at $\phi_1(\hat{p}_k) = \hat{\phi}_{k-1}$, which is equal to expression (12) (note that $(*)$ in expression (13) is equal to $\hat{C}_{k-1}$ when $\phi_1(\hat{p}_k) = \hat{\phi}_{k-1}$).

Part 2.2. We now argue that, in any equilibrium and for any $\phi_0 > \hat{\phi}_k$, a payoff arbitrarily close to expression (12) can be achieved by charging a price slightly below $\hat{p}_k$. Indeed, assume that the seller offers $\hat{p}_k - \varepsilon$, for a small $\varepsilon > 0$. If the h-buyer rejects the offer for sure, then $\phi_1(\hat{p}_k - \varepsilon) > \hat{\phi}_k$. Still, the continuation payoff of the h-buyer from rejecting is at most $\delta_{\text{b}}\, (h - \hat{p}_k)$ (by the induction hypothesis), which is smaller than $h - (\hat{p}_k - \varepsilon)$ (i.e., the payoff of accepting $\hat{p}_k - \varepsilon$ in the first period) if ε is small enough, contradicting the incentive to reject the first offer. Alternatively, the h-buyer cannot be strictly willing to accept $\hat{p}_k - \varepsilon$, since otherwise, the next period's price is ℓ, making rejection a strictly profitable deviation. Hence, the h-buyer must be indifferent between accepting and rejecting $\hat{p}_k - \varepsilon$, and so his continuation payoff in the second period should be $\delta_{\text{b}}^{-1}\, (h - (\hat{p}_k - \varepsilon))$. From the definition of $\hat{p}_k$ we have that, if ε is small enough, such a continuation payoff is in $(h - \hat{p}_{k-1}, h - \hat{p}_{k-2})$, so it must be that $\phi_1(\hat{p}_k - \varepsilon) = \hat{\phi}_{k-1}$.

An implication of the previous argument is that, since we had argued in Part 2.1 that (12) is an upper bound on the seller's payoff when $\phi_0 \in [\hat{\phi}_k, \hat{\phi}_{k+1})$, and now we obtained that it is also a lower bound, we have that (12) is the unique equilibrium seller's payoff in this range of prior beliefs. It is only left to obtain the value of $\hat{\phi}_{k+1}$.

Part 2.3. Define $\hat{p}_{k+1} := (1 - \delta_{\text{b}})\, h + \delta_{\text{b}}\, \hat{p}_k$. Note that if $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$ and the seller offers a price $p_0 \in (\hat{p}_k, \hat{p}_{k+1})$ then, necessarily, it must be that $\phi_1 = \hat{\phi}_k$. Indeed, if $\phi_1 > \hat{\phi}_k$ then the continuation value of the h-buyer is at most $\delta_{\text{b}}\, (h - \hat{p}_k) < h - p_0$, contradicting that $\phi_1 > \hat{\phi}_k$; while if $\phi_1 < \hat{\phi}_k$ then the continuation value of the h-buyer is larger than $\delta_{\text{b}}\, (h - \hat{p}_k) > h - p_0$, contradicting that $\phi_1 < \hat{\phi}_k$. This implies that a lower bound on the seller's equilibrium payoff when $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$

is

$$\frac{\phi_0-\hat{\phi}_k}{1-\hat{\phi}_k}\,\hat{p}_{k+1}+\frac{1-\phi_0}{1-\hat{\phi}_k}\,\delta_{\mathrm{s}}\,\hat{C}_k\ . \tag{14}$$

The previous expression increases in ϕ_0 faster than expression (12). Indeed, both expressions are linear in ϕ_0, expression (14) is higher than expression (12) when $\phi_0=1$, and expression (14) is lower than expression (12) when $\phi_0=\hat{\phi}_k$.[16] It is then the case that offering $\hat{p}_k$ in period 1 is optimal only if $\phi_0\leq\hat{\phi}'_{k+1}$, where $\hat{\phi}'_{k+1}>\hat{\phi}_k$ is the unique ϕ_0 which makes expression (14) equal to (12). Hence, we have $\hat{\phi}_{k+1}\leq\hat{\phi}'_{k+1}$.

Part 2.4. Similar to Part 1 of this proof, define

$$\overline{C}^{\varepsilon}_{k+1}:=\sup\left(\cup_{\phi_0\in[\hat{\phi}_{k+1},\hat{\phi}_{k+1}+\varepsilon]}\,\mathcal{C}^*(\phi_0|\,\Pr(p_0>\hat{p}_k)>0\right)$$

for all $\varepsilon>0$, and we let $\overline{C}^{0}_{k+1}:=\lim_{\varepsilon\searrow 0}\overline{C}^{\varepsilon}_{k+1}$. We let $(\phi_0^n)_{n=1}^{\infty}$ be a decreasing sequence converging to $\hat{\phi}_{k+1}$ with a corresponding sequence $(\pi^n,\alpha^n,\phi^n)_{n=1}^{\infty}$ satisfying that (i) for all n, (π^n,α^n,ϕ^n) is an equilibrium when the prior is ϕ_0^n satisfying $\pi^n(p_0^n>\hat{p}_k)>0$, and (ii) $\lim_{n\to\infty}C_0^n=\overline{C}^0_{k+1}$. Note that, without loss of generality for the argument, we can assume that the seller plays a pure strategy (i.e., $\pi^n(p_0^n)=1$ for some $p_0^n>\hat{p}_k$) for all n (see Footnote 15). We now consider five cases, which are analogous to the four cases in Part 1 with one extra case:

1. Assume, for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $\pi^{n_m}(p_1^{n_m}>\hat{p}_k|p_0^{n_m})>0$ for all m. By the induction hypothesis, this implies that $\phi_1^{n_m}(p_0^{n_m})\in[\hat{\phi}_{k+1},\phi_0^{n_m}]$. The payoff of the seller is

$$C_0^{n_m}=\frac{\phi_0^{n_m}-\phi_1^{n_m}(p_0^{n_m})}{1-\phi_1^{n_m}(p_0^{n_m})}\,p_0^{n_m}+\frac{1-\phi_0^{n_m}}{1-\phi_1^{n_m}(p_0^{n_m})}\,\delta_{\mathrm{s}}\,C_1^{n_m}(p_0^{n_m})\ .$$

By assumption, the left-hand side tends to $\overline{C}^0_{k+1}$ as $m\to\infty$. The first term on the right-hand side tends to 0 because $\phi_0^{n_m}\to\hat{\phi}_{k+1}$ as $m\to\infty$, hence $\phi_1^{n_m}(p_0^{n_m})\to\hat{\phi}_{k+1}$ as $m\to\infty$ as well. Since $C_1^{n_m}(p_0^{n_m})\leq\overline{C}^{\varepsilon}_{k+1}$ for $\varepsilon=\phi_1^{n_m}(p_0^{n_m})-\hat{\phi}_{k+1}$, we have that

$$\overline{C}^0_{k+1}=\lim_{m\to\infty}C_0^{n_m}=\delta_{\mathrm{s}}\lim_{m\to\infty}C_1^{n_m}(p_0^{n_m})\leq\delta_{\mathrm{s}}\,\overline{C}^0_{k+1}\ ;$$

hence, $\overline{C}^0_{k+1}\leq 0$. This is a contradiction because $\overline{C}^0_{k+1}\geq\ell>0$.

[16]Note that, for $\phi_0=\hat{\phi}_k$, (12) is equal to $\hat{C}_k$ and (14) is equal to $\delta_{\mathrm{s}}\,\hat{C}_k<\hat{C}_k$.

2. Assume, again for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $\pi^n(p_1^{n_m} < \hat{p}_k | p_0^{n_m}) > 0$ for all m, which implies that $\phi_1^{n_m}(p_0^{n_m}) \le \hat{\phi}_k$ and that the continuation value of the h-buyer in the second period is higher than $h - \hat{p}_k$. This implies that the h-buyer rejects $p_0^{n_m}$ for sure, hence $\phi_1^{n_m}(p_0^{n_m}) = \phi_0^{n_m} > \hat{\phi}_k$, a contradiction.

3. Assume now, again for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $p_0^{n_m} > \hat{p}_{k+1}$ for all m. Since, from the previous result, we have that $\pi^{n_m}(p_1^{n_m} = \hat{p}_k | p_0^{n_m}) = 1$ if m is large enough, we have that the h-buyer accepts with probability zero the first price. Now, we have $C_0^{n_m} = \delta_{\text{s}} C_1^{n_m}(p_0^{n_m})$, which again implies that $\overline{C}_{k+1}^0 \le \delta_{\text{s}} \overline{C}_{k+1}^0$, a contradiction.

4. Assume, again for the sake of contradiction, that there is a strictly increasing sequence $(n_m)_{m=1}^{\infty}$ such that $p_0^{n_m} \in (\hat{p}_k, \hat{p}_{k+1})$ for all m. Since, from the previous results, we have $\pi^{n_m}(p_1^{n_m} = \hat{p}_k | p_0^{n_m}) = 1$, the h-buyer accepts such price for sure, implying again that $p_1^{n_m}(p_0^{n_m}) = \ell$, and so that the h-buyer is strictly willing to reject $p_0^{n_m}$, a contradiction.

5. The only possibility left is that, if n is large enough, then $p_0^n = \hat{p}_{k+1}$ and $\pi^n(p_1^n = \hat{p}_k | p_0^n) = 1$. Hence, we have that $\overline{C}_{k+1}^0$ equal to (14). Since, as we argued in Part 2.2, $\overline{C}_{k+1}^0$ is weakly higher than (12), we have that $\hat{\phi}_{k+1} \ge \hat{\phi}'_{k+1}$.

Overall, we conclude that $\hat{\phi}_{k+1} = \hat{\phi}'_{k+1}$. Hence, in any equilibrium, if $\phi_0 \in (\hat{\phi}_k, \hat{\phi}_{k+1})$, the seller offers $\hat{p}_k$ for sure in the first period (and the continuation play is according to the statement of the induction hypothesis), while if $\phi_0 = \hat{\phi}_{k+1}$, the seller potentially mixes between $\hat{p}_k$ and $\hat{p}_{k+1}$ (and also the continuation play is according to the statement of the induction hypothesis). Since we argued that offering a price below $\hat{p}_{k+1}$ is dominated by offering a price slightly below $\hat{p}_{k+1}$ if $\phi_0 > \hat{\phi}_{k+1}$, which proves the second point of the induction hypothesis.

Existence of an equilibrium. The proof of existence of an equilibrium is by construction. To do that, for all $p \in [\ell, h)$ and $\phi \in [0,1)$, we let $k(p)$ and $k(\phi)$ indicate the unique values of such that $p \in [\hat{p}_{k(p)}, \hat{p}_{k(p)+1})$ and $\phi \in [\hat{\phi}_{k(\phi)}, \hat{\phi}_{k(\phi)+1})$, respectively. We then recursively define, for each history $p^t = (p_0, ..., p_{t-1})$ and price p_t,

$$\phi(p_t, p^t) := \begin{cases} \phi(p^t) & \text{if } p_t \ge h, \\ \min\{\hat{\phi}_{k(p_t)}, \phi(p^t)\} & \text{if } p_t \in [\ell, h), \\ 0 & \text{if } p_t < \ell, \end{cases}$$

and

$$\pi(\cdot|p^t) = \begin{cases} \beta(p_{t-1}) \circ \hat{p}_{k(p_{t-1})} + (1-\beta(p_{t-1})) \circ \hat{p}_{k(p_{t-1})-1} & \text{if } p_{t-1} \geq \ell \text{ and } \phi(p^t) \geq \hat{\phi}_{k(p_{t-1})}, \\ 1 \circ \hat{p}_{k(\phi(p^t))} & \text{otherwise,} \end{cases}$$

where $\beta(p_{t-1}) := \frac{p_{t-1} - \hat{p}_{k(p_{t-1})}}{(1-\delta_b)(h - \hat{p}_{k(p_{t-1})})}$. For example, if $\phi(p^t) \in (\hat{\phi}_1, \hat{\phi}_2)$ and $p_t \in [\hat{p}_1, \hat{p}_2]$, then the buyer accepts so that $\phi(p_t, p^t) = \hat{\phi}_1$, and in period $t+1$ the seller randomizes between ℓ and $\hat{p}_1$ so that the h-buyer is indifferent between accepting or not p_t a time t. Alternatively, if $\phi(p^t) \in (\hat{\phi}_1, \hat{\phi}_2)$ and $p_t > \hat{p}_2$, the buyer rejects for sure, while if $p_t < \hat{p}_1$, then the h-buyer accepts for sure and the ℓ-buyer accepts if $p_t \leq \ell$. It is easy to see that (ϕ, π) defined above, together with the implied buyer strategy α, form a perfect Bayesian equilibrium. □

Proofs of Corollaries 3.1-3.3

Proof. The proofs follow from the arguments in the main text. □

Proof of Theorem 3.2

Proof. We now look for seller strategies π and maps from each seller's strategy π' to a strategy of the buyer $(\alpha_\ell(\cdot|\pi'), \alpha_h(\cdot|\pi'))$ such that (i) π maximizes the $C(\varnothing|\pi', \alpha_\ell, \alpha_h, \phi_0)$ among all π' and (ii) each $\alpha_\theta(\cdot|\pi')$ maximizes $V_\theta(p^t; \pi', \alpha_\theta)$ for all θ, p^t, and π' (note that, differently from Definition 2.1, we do not require the seller's strategy to be sequentially optimal and we allow the buyers to observe the strategy of the seller). It is not difficult to see that we can assume, without loss of generality, that the buyer purchases when he is indifferent.[17] We then assign, to each seller's strategy π, the payoff $C_0(\pi)$ computed under the assumption that each type of the buyer buys in the first period it is optimal for him to do so. We look for π maximizing $C_0(\pi)$.

We first argue that it is without loss of optimality to focus on equilibria where the seller's price in the first period is non-stochastic. To see that, pose a seller strategy π and let p_0 be a price in the support of the seller's offer distribution in the first period such that, conditionally on offering p_0 in the first period, the seller obtains a payoff $\hat{C}_0 \geq C_0(\pi)$. Note that there exists a seller strategy where the seller offers p_0 for sure in the first period and the corresponding

[17]The argument is standard: It is easy to see that, for a given π, the seller prefers each type of the buyer to buy at the earliest time which is optimal for the buyer, and that she can slightly alter her strategy to make such acceptance strictly optimal.

continuation price path afterward, which gives a payoff equal to $\hat{C}_0$ to the seller.

From the previous observation, it follows that if there is a seller's strategy where the h-buyer does not purchase at time 0 for sure, there is another seller's strategy where the h-buyer purchases at time 0 for sure, which gives the seller a higher payoff. Hence, we focus without loss of generality on seller strategies the h-buyer purchases at time 0 for sure.

Now, see that if a seller strategy is such that the transaction price is strictly below ℓ with positive probability, there is another seller strategy where no price is strictly below ℓ, which gives the seller a higher payoff. Such price path can be obtained by replacing each instance where a price $p_t < \ell$ is offered by a price equal to $\beta p_t + (1-\beta)\ell$ for some $\beta \in (0,1)$. It is easy to see that the h-buyer still prefers to buy at time 0, while the ℓ-buyer buys at a weakly earlier time at a weakly higher price; hence, the seller is weakly better off.

We now study stochastic price paths that maximize the seller's payoff conditional on a price $p_0 \in \{\ell\} \cup [\hat{p}_1, h]$ being offered in period 0 and accepted for sure by the h-buyer, while the ℓ-seller buys at the first (random) time $\tilde{t}$ where the price is ℓ.[18] It is clear that the seller's optimality requires that the h-buyer is indifferent between buying at p_0 and mimicking the ℓ-buyer. Hence, it must be that

$$h - p_0 = \mathbb{E}[\delta_{\mathrm{b}}^{\tilde{t}}]\,(h-\ell)\ .$$

Assume that $\tilde{t}$ is optimal and assigns positive probability to some time $t_1 > 1$. The seller can replace the event where ℓ is offered at t_1 by a lottery between $t_1 - 1$ (with probability q) and $t_1 + 1$ (with probability $1-q$) so that the h-buyer remains indifferent, that is, such that

$$\delta_{\mathrm{b}}^{t_1} = q\,\delta_{\mathrm{b}}^{t_1-1} + (1-q)\,\delta_{\mathrm{b}}^{t_1+1} \quad \Rightarrow \quad q = \tfrac{\delta_{\mathrm{b}}}{1+\delta_{\mathrm{b}}}\ .$$

The change in the seller's payoff conditional on this even occurring is

$$\left(q\,\delta_{\mathrm{s}}^{t_1-1} + (1-q)\,\delta_{\mathrm{s}}^{t_1+1}\right)\ell - \delta_{\mathrm{s}}^{t_1}\,\ell = \tfrac{1-\delta_{\mathrm{s}}}{1+\delta_{\mathrm{b}}}\,\delta_{\mathrm{s}}^{t_1-1}\,(\delta_{\mathrm{b}} - \delta_{\mathrm{s}})\,\ell\ .$$

There are two cases:

[18] Note that if $p_0 \in (\ell, \hat{p}_1)$, there is no strategy of the seller that makes the h-buyer indifferent between buying in period 0 or buying at a later date.

1. Assume first $\delta_b \geq \delta_s$.[19] In this case, the seller weakly benefits from the previous change. An analogous argument to that before shows that, without loss, we can assume that the support of $\tilde{t}$ in an optimal seller strategy is either $\{0\}$ or a subset of $\{1, +\infty\}$ in this case. That is, either the seller offers ℓ in period 0, or offers p_0 in period 0 and then price ℓ in period 1 with probability $q_1 \in [0,1]$ satisfying

$$h - p_0 = \delta_b \, q_1 \, (h - \ell) \ .$$

Then, the seller's payoff is

$$\max \left\{ \ell, \max_{q_1 \in [0,1]} \left(\phi_0 \, (h - \delta_b \, q_1 \, (h - \ell)) + (1 - \phi_0) \, \delta_s \, q_1 \, \ell \right) \right\} \ .$$

It is then easy to see that, in the first period and under any optimal seller strategy, the seller either offers ℓ (if $\phi_0 < \phi^*$) or h (if $\phi_0 > \phi^*$) or randomizes between the two (if $\phi_0 = \phi^*$), and that no trade takes place for all $t > 0$.

2. If $\delta_b < \delta_s$, the seller loses from the change described above. The implication now is that the support of $\tilde{t}$ in any seller strategy is either $\{0\}$ or a subset of $\{\bar{t}, \bar{t}+1\}$ for some $\bar{t} > 0$. It is clear that, under an optimal seller strategy, offering ℓ in the first period is optimal for the seller if $\phi_0 < \hat{\phi}_1$, while offering some $p_0 > \ell$ in the first period is optimal if $\phi_0 > \hat{\phi}_1$. Assume the second case, and let $q \in [0,1]$ be such that the seller offers price ℓ with probability q at $\bar{t}$ and with probability $1 - q$ at $\bar{t}+1$. The seller maximizes

$$\phi_0 \, p_0 + (1 - \phi_0) \, (q \, \delta_s^{\bar{t}} + (1 - q) \, \delta_s^{\bar{t}+1}) \, \ell$$

over $\bar{t}$ and q subject to

$$h - p_0 = (q \, \delta_b^{\bar{t}} + (1 - q) \, \delta_b^{\bar{t}+1}) \, (h - \ell) \ .$$

Standard analysis implies that the optimal $\bar{t}$ is the smallest satisfying

$$(1 - \phi_0) \, (1 - \delta_s) \, \delta_s^{\bar{t}} \, \ell - \phi_0 \, (1 - \delta_b) \, \delta_b^{\bar{t}} \, (h - \ell) \geq 0 \tag{15}$$

[19] The main text describes a shorter argument using the result in Stokey (1979). We provide here a self-contained proof.

while the optimal q is 1 if the previous expression holds with strict inequality for the optimal $\bar{t}$, or any value in $[0,1]$ if the previous expression holds with equality. The statement of the Theorem 3.2 follows from these observations. Note that, since equation (15) holds with equality for some $\bar{t} \in \{1,2,...\}$ non-generically in the parameters of the model,[20] we have that, generically, $q=1$. □

Proof of Theorem 4.1

Proof. Like the proof of Theorem 3.1, the proof proceeds by induction, this time over the length of the horizon $T=0,1,2,...$ The following will be our induction hypothesis:

Induction hypothesis for T: There exists an increasing sequence $(k_{T'},\overline{\phi}_{T'})_{T'=0}^{T}$, with $\overline{\phi}_{T'} \geq \hat{\phi}_{k_{T'}}$ for all T' and $\overline{\phi}_0 = \phi^* := \ell/h$, such that, in any equilibrium, the following holds true:

1. If $\phi_0 > \overline{\phi}_T$ then the on-path history is $(h,h,...,h)$ and the corresponding beliefs are $(\phi_0,\overline{\phi}_{T-1},...,\overline{\phi}_0)$.
2. If $\phi_0 < \overline{\phi}_T$ then,
   (a) if $\phi_0 \in (\hat{\phi}_{k_T},\overline{\phi}_T)$ then the on-path history is $(\hat{p}_{k_T},\hat{p}_{k_T-1},...,\ell)$ and the corresponding beliefs are $(\phi_0,\hat{\phi}_{k_T-1},...,\hat{\phi}_0)$, and,
   (b) otherwise, the equilibrium is as specified in Theorem 3.1.
3. If $\phi_0 = \overline{\phi}_T$ then the seller randomizes between part 1 and part 2.[21]

Part 1: Proof for $T=0$ and $T=1$. The result is clear for $T=0$, where $\overline{\phi}_0 = \phi^*$ and $k_0 = 0$. The result for $T=1$ follows from the argument in the main text.

Part 2: Proof for $T+1>1$. Fix some $T \geq 1$ and assume that the induction hypothesis holds for T. By the same argument as in the proof of Theorem 3.1, in any equilibrium, if the seller offers (on or off-path) a price $p_0 \in (\hat{p}_k,\hat{p}_{k+1})$ for some $k \leq k_T$ such that $\phi_0 \geq \hat{\phi}_k$, then the h-buyer must be indifferent between accepting it or not; so $\phi_1 = \hat{\phi}_k$ and the seller obtains

$$\tfrac{\phi_0-\phi_1}{1-\phi_1}\, p_0 + \tfrac{1-\phi_0}{1-\phi_1}\, \delta_{\mathrm{s}}\, \hat{C}_k \ . \tag{16}$$

[20]That is, equation (15) holds with equality only if $\log\left(\frac{\phi_0}{1-\phi_0}\,\frac{1-\delta_{\mathrm{b}}}{1-\delta_{\mathrm{s}}}\,\frac{h-\ell}{\ell}\right)/\log(\delta_{\mathrm{s}}/\delta_{\mathrm{b}})$ is a natural number.

[21]That is, if $\phi_0 = \overline{\phi}_T > \hat{\phi}_{k_T}$ then the seller randomizes between the paths described in parts 1 and 2(a), and if $\phi_0 = \overline{\phi}_T = \hat{\phi}_{k_T}$ then the seller randomizes between the paths described in parts 1 and 2(b).

If, instead, the seller offers $\hat{p}_{k+1}$ for some $k \leq k_T$ such that $\phi_0 \geq \hat{\phi}_k$, then $\phi_1 \in [\hat{\phi}_k, \min\{\phi_0, \hat{\phi}_{k+1}, \overline{\phi}_T\}]$, and the payoff of the seller is again given by (16). If $\phi_0 \geq \overline{\phi}_T$ and $p_0 \in (\hat{p}_{k_T+1}, h)$, then $\phi_1 = \overline{\phi}_T$ and the seller's payoff is

$$\overline{C}_{T+1}(\phi_0, p_0) := \tfrac{\phi_0 - \overline{\phi}_T}{1-\overline{\phi}_T}\, p_0 + \tfrac{1-\phi_0}{1-\overline{\phi}_T}\, \delta_{\mathrm{s}}\, \overline{C}_T \ , \tag{17}$$

where $\overline{C}_T$ is the continuation payoff at T when the posterior is $\overline{\phi}_T$. Finally, if $\phi_0 \geq \overline{\phi}_T$ and $p_0 = h$, then $\phi_1 \in [\overline{\phi}_T, \phi_0]$ and the payoff of the seller is

$$\tfrac{\phi_0 - \phi_1}{1-\phi_1}\, h + \tfrac{1-\phi_0}{1-\phi_1}\, \delta_{\mathrm{s}} \left(\tfrac{\phi_1 - \overline{\phi}_T}{1-\overline{\phi}_T}\, h + \tfrac{1-\phi_1}{1-\overline{\phi}_T}\, \overline{C}_T \right) . \tag{18}$$

It is easy to see that both (16) and (18) are increasing in ϕ_1, and that expression (16) is smaller than (18) for p_0 close enough to h.

Because, in the T-period model, the seller is indifferent between offering h and $\hat{p}_{k_T}$ when $\phi_0 = \overline{\phi}_T$, we have

$$\overline{C}_T = \tfrac{\overline{\phi}_T - \hat{\phi}_{k_T-1}}{1-\hat{\phi}_{k_T-1}}\, \hat{p}_k + \tfrac{1-\overline{\phi}_T}{1-\hat{\phi}_{k_T-1}}\, \delta_{\mathrm{s}}\, \hat{C}_{k_T-1} \ .$$

Hence, for ϕ_0 slightly below $\overline{\phi}_T$, the payoff of offering a price slightly below h in the $(T+1)$-period model is approximately $\delta_{\mathrm{s}} \overline{C}_T$. The payoff from offering a price slightly below $\hat{p}_k$ is approximately equal to $\overline{C}_T$. It is then clear that for all $\phi_0 \leq \overline{\phi}_T$, the price offered in the first period in the $(T+1)$-model is that given by the induction hypothesis.

Assume $\phi_0 \geq \overline{\phi}_T$. In this case, by offering a price slightly below h, the seller can obtain a payoff arbitrarily close to $\overline{C}_{T+1}(\phi_0, h)$ defined in equation (17). It is then easy to see that k_{T+1} and $\overline{\phi}_{T+1}$ are determined as follows:

1. If $\overline{C}_{T+1}(\overline{\phi}_T, h) = \hat{C}_{k_T}$ then $k_{T+1} := k_T + 1$ and $\overline{\phi}_{T+1} := \hat{\phi}_{k_T+1}$.
2. If $\overline{C}_{T+1}(\overline{\phi}_T, h) < \hat{C}_{k_T}$ then $k_{T+1} := k_T + 1$ and $\overline{\phi}_{T+1}$ is the unique solution to

$$\overline{C}_{T+1}(\overline{\phi}_{T+1}, h) = \tfrac{\overline{\phi}_{T+1} - \hat{\phi}_{k_T}}{1-\hat{\phi}_{k_T}}\, \hat{p}_{k_T+1} + \tfrac{\overline{\phi}_{T+1} - \hat{\phi}_{k_T}}{1-\hat{\phi}_{k_T}}\, \delta_{\mathrm{s}}\, \hat{C}_{k_T-1} \ .$$

3. If $\overline{C}_{T+1}(\overline{\phi}_T, h) > \hat{C}_{k_T}$ then $k_{T+1} := k_T$ and $\overline{\phi}_{T+1}$ is the unique solution to

$$\overline{C}_{T+1}(\overline{\phi}_{T+1}, h) = \tfrac{\overline{\phi}_{T+1} - \hat{\phi}_{k_T-1}}{1-\hat{\phi}_{k_T-1}}\, \hat{p}_{k_T} + \tfrac{\overline{\phi}_{T+1} - \hat{\phi}_{k_T-1}}{1-\hat{\phi}_{k_T-1}}\, \delta_s\, \hat{C}_{k_T-2}\ .$$

Part 2 of this proof is then concluded.

Existence of an equilibrium. A perfect Bayesian equilibrium can be constructed similarly as the one constructed in the proof of Theorem 3.1. □

Proof of Corollary 4.1

Proof. The proof follows from a formalization of the arguments in the main text. □

References

Abreu, D. and F. Gul (2000): "Bargaining and reputation," *Econometrica*, 68, 85–117.

Ausubel, L. M. and R. J. Deneckere (1989): "Reputation in bargaining and durable goods monopoly," *Econometrica*, 57, 511–531.

Bagnoli, M., S. W. Salant, and J. E. Swierzbinski (1989): "Durable-goods monopoly with discrete demand," *Journal of Political Economy*, 97, 1459–1478.

Board, S. and M. Pycia (2014): "Outside options and the failure of the Coase conjecture," *American Economic Review*, 104, 656–671.

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

Daley, B. and B. Green (2020): "Bargaining and news," *American Economic Review*, 110, 428–74.

Deneckere, R. and M.-Y. Liang (2006): "Bargaining with interdependent values," *Econometrica*, 74, 1309–1364.

Dilmé, F. (2023a): "Repeated bargaining with imperfect information about previous transactions," .

——— (2023b): "The role of discounting in bargaining with private information," .

Dilmé, F. and F. Li (2019): "Revenue management without commitment: Dynamic pricing and periodic flash sales," *The Review of Economic Studies*, 86, 1999–2034.

Fuchs, W. and A. Skrzypacz (2010): "Bargaining with arrival of new traders," *American Economic Review*, 100, 802–836.

——— (2013): "Bargaining with deadlines and private information," *American Economic Journal: Microeconomics*, 5, 219–43.

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

Gerardi, D., L. Maestri, and I. Monzon (2022): "Bargaining over a Divisible Good in the Market for Lemons," *American Economic Review*, 112, 1591–1620.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture," *Journal of Economic Theory*, 39, 155–190.

Güth, W. and K. Ritzberger (1998): "On durable goods monopolies and the Coase-conjecture," *Review of Economic Design*, 3, 215–236.

Kaya, A. and K. Kim (2018): "Trading dynamics with private buyer signals in the market for lemons," *The Review of Economic Studies*, 85, 2318–2352.

Kaya, A. and S. Roy (2022): "Repeated Trading: Transparency and Market Structure," .

Landsberger, M. and I. Meilijson (1985): "Intertemporal price discrimination and sales strategy under incomplete information," *RAND Journal of Economics*, 16, 424–430.

Lauermann, S. and A. Wolinsky (2016): "Search with adverse selection," *Econometrica*, 84, 243–315.

McAfee, R. P. and T. Wiseman (2008): "Capacity choice counters the Coase conjecture," *The Review of Economic Studies*, 75, 317–331.

Nava, F. and P. Schiraldi (2019): "Differentiated durable goods monopoly: a robust Coase conjecture," *American Economic Review*, 109, 1930–1968.

Sobel, J. (1991): "Durable goods monopoly with entry of new consumers," *Econometrica*, 59, 1455–1485.

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