

Repeated Trade with Imperfect Information about Previous Transactions

Francesc Dilmé

9th August 2024

This paper studies repeated trade with noisy information about previous transactions. A buyer has private information about his willingness to pay, which is either low or high, and buys goods from different sellers over time. Each seller observes a noisy history of signals about the buyer's previous purchases and sets a price. We compare the cases where previous prices are observable to sellers with the case where they are not. We show that, in both cases, more signal precision is counterbalanced in equilibrium by less separation between types, either because (1) sellers offer discounted prices more often, which are accepted by both types, or (2) the buyer with a high valuation rejects high prices with a higher probability. The effect on welfare and efficiency of making prices observable depends on the signal precision: When the signal is imprecise, making prices public strengthens the discounting mechanism, improving efficiency and buyer welfare; when the signal is precise, making prices public activates the rejection mechanism, and efficiency and buyer welfare may decrease. Independently of price observability, the buyer tends to benefit from a more precise signal about previous purchases.

Keywords: Repeated trade, asymmetric information, internet cookies

JEL codes: C73, C78, D82.

Department of Economics, University of Bonn, fdilme@uni-bonn.de. I thank Sarah Auster, Doruk Cetemen, and Stephan Lauer mann, and the participants in the Bonn Theory Workshop and City Economic Theory Conference for their helpful comments. 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 many markets, participants conduct multiple transactions over time. For instance, people shopping online may make multiple purchases from various retailers, and companies and government entities procure goods and services from multiple suppliers. In these markets, some information regarding past transactions is often accessible to traders; for example, internet cookies record prior consumer interactions, public reports document the terms of procurement contracts, and information may be disseminated through leaks, word of mouth, or online forums. Hence, a trader's behavior is affected by the fact that other traders may acquire information about his decisions and use it to their advantage.

The current theoretical understanding of how information about previous transactions affects decentralized pricing is limited and based on two extreme cases: that of perfect information, and that of no information. In the first case, Hart and Tirole (1988) show that Coasian forces determine the equilibrium outcome. The basic logic is clear. If a buyer with high willingness to buy accepts a high price now, he obtains a small surplus from the transaction but also signals his high willingness to buy, which weakens his bargaining position in the future. Thus, in equilibrium he rejects high prices; this means he loses a small surplus now but secures better prices in the future. To avoid rejection, sellers offer him low prices. In the second case, when there is no information about previous transactions, the buyer's signaling concern disappears, and the repeated static outcome results in equilibrium.

In practice, these extreme cases are rare, whereas intermediate cases with imperfect and noisy information are common. Current models based on the extremes, therefore, give little guidance on the possible implications of regulatory or technological changes affecting the information available to traders.

This paper's first contribution is to provide a tractable model to study how noisy information about previous transactions affects repeated trade. A long-lived buyer meets a short-lived seller at each instant in continuous time. The buyer has a permanent type, corresponding to his valuation of the seller's good, which is either low (ℓ) or high (h), with $0 < \ell < h$. (The buyer's type may also be interpreted as his wealth or marginal value for money.) Each seller sets a uniform price, and the buyer decides the amount of the good he buys. We initially assume that each seller observes the previous price offers, together with a history of noisy signals about the buyer's previous acceptance decisions. We study Markov perfect equilibria with the sellers' posterior belief about the high type as the Markov variable, and with ℓ always accepted.

We find that imperfect information on acceptance decisions averts Coasian dynamics while generating non-trivial trade dynamics. Intuitively, if ℓ were offered and accepted for sure in an interval of high posteriors, the buyer's continuation value would be flat in this interval. A seller could then benefit from offering a price slightly below h : such a price would be accepted by the high-valuation buyer (the

h -buyer), because it would give him a positive surplus while not significantly changing the sellers' beliefs about his type. In the unique equilibrium, while only low prices are offered for low posteriors, high prices are offered with positive probability for intermediate posteriors and for sure for high posteriors. Two incentive constraints are then satisfied for intermediate and high posteriors: because a seller can ensure trade by offering ℓ , he offers a high price in equilibrium only if the h -buyer accepts it with high enough probability, while the h -buyer accepts a high price only if doing so does not decrease his continuation payoff too much.

The paper's second main contribution is to present two equilibrium mechanisms that keep the incentives of the buyer and sellers in balance. The first is a *discount mechanism* for intermediate-low posteriors: sellers offer the discounted price ℓ with positive probability. Such discounted prices slow learning (because the acceptance signal is uninformative when price ℓ is offered, as both buyer types accept it) and increase the continuation value of the h -buyer, making it less sensitive to the arrival of information. As a result, the h -buyer is willing to accept high prices in equilibrium, which means that sellers will offer high prices even when the probability of the buyer's having a low valuation (and hence rejecting the price) is relatively high. The second mechanism is a *rejection mechanism*, which arises for intermediate-high posteriors when the signal is sufficiently informative; in this case, the h -buyer rejects high prices with positive probability. This mechanism further slows learning, because a signal indicating rejection of a high price does not necessarily indicate that the buyer's valuation is low.

Our mechanisms correspond to the two possible equilibrium mechanisms for lowering the informativeness of acceptance decisions as required to balance the incentives in the dynamic model: the first entails a very high acceptance probability (sellers make low price offers that are likely accepted), and the second entails a very low acceptance probability (sellers make high price offers that are likely rejected).¹ In either case, a contrarian signal is interpreted as likely noise, while a conformist signal is interpreted as carrying little information. We believe that these mechanisms extend beyond the particular assumptions of our model.

The paper's third main contribution is its analysis of the welfare effects of policies affecting the precision of the information available to sellers—for example, internet privacy regulations or transparency laws for firms and government agencies. For the extremes of signal informativeness, we recover the existing results: the price tends to ℓ as the signal becomes arbitrarily precise and to the static monopolistic price as the signal tends toward pure noise. When the signal precision is low, reducing it weakens the discount mechanism, which decreases equilibrium trade and increases equilibrium prices; hence, buyer welfare and efficiency decrease. When the signal precision is high, reducing it weakens both the discount

¹ Note that, if α is the unconditional probability that a given price is accepted, both the variance and the entropy of the implied Bernoulli random variable are single-peaked in α .

and rejection mechanisms. Because weakening the rejection mechanism increases trade, the total effect on efficiency and buyer welfare is unclear. We prove that there is a large region of intermediate posteriors in which the total effect of reducing the signal precision is detrimental to efficiency and buyer welfare.²

We also analyze how the effects of signal precision on welfare depend on the observability of prices. For this we consider a modification of our main model in which sellers observe the previous acceptance signals but not the prices offered. In this model, deviations by sellers are unobservable, which implies that all equilibrium offers are accepted for sure; thus, the rejection mechanism is absent. Nevertheless, we show that when the signal precision is low, buyer welfare and efficiency are lower than in the model where prices are observable to sellers. The reason is that the equilibrium information carried by the signal is higher if sellers know the prices offered to the buyer in past transactions: when prices are observable, the buyer's acceptance of a high price is very informative; when prices are unobservable, sellers cannot distinguish between acceptance of high prices and acceptance of low prices. Thus, when prices are observable, a stronger discounting mechanism is needed to maintain the h -buyer's incentive to accept high prices, and buyer welfare and efficiency are higher as a result. However, when the signal precision is high, we show that this result may be reversed: observability of prices may be detrimental to the buyer. With high signal precision, when prices are observable, the rejection mechanism keeps the posterior high for a longer time, lowering the value to the h -buyer of mimicking the ℓ -buyer. Furthermore, observability of prices may lower market efficiency in this case, by causing the h -buyer to trade less often.

Technical contribution: The paper provides a tractable continuous-time framework for analyzing repeated decentralized pricing in a setting where uninformed traders have all the bargaining power and where private information is revealed through acceptance of price offers. Our model differs from the continuous-time models in the previous literature in that we modify the recursive formulation of the problem to allow both the buyer and the sellers to use instantaneous randomizations over non-trivial regions of posteriors. Such mixed strategies play a crucial role in encoding the various economic mechanisms through which the endogenous pricing and acceptance decisions affect learning in the unique equilibrium. We are also able to accommodate and compare within the same setting the two cases that are often studied in discrete-time trade models: the public-offers case and the private-offers case.

Regulatory implications: Our paper provides a novel perspective on the study of pricing in markets with repeated trade. In our model, changes in signal precision affect the equilibrium behavior of informed and

² Note that the unobservability of prices affects equilibrium dynamics in two important ways. First, because sellers randomize in a region of the belief space, future sellers are unsure about the meaning of the signal when prices are unobservable. Second, even in regions of the belief space where sellers do not randomize, the incentives of the sellers and the buyer differ depending on the observability of prices: When prices are observable, the signal informativeness depends on both the price offered and the acceptance decision, while it only depends on the acceptance decision when prices are unobservable. These considerations affect the incentives of the sellers and the buyer to deviate and shape the equilibrium strategies.

uninformed traders, affecting welfare and efficiency in non-trivial ways. In Section 5, we discuss some policy implications in light of current debates on privacy and transparency in online or procurement markets. While we focus on personalized pricing, we believe our setting captures the main forces in play in other mechanisms used in practice, such as personalized advertising or recommendations.

1.1 Literature review

As explained above, Hart and Tirole (1988) show that in a model with perfect information about the previous transactions, Coasian forces favor an equilibrium in which prices are equal to the buyer's lowest valuation, as the acceptance of a high price results in a permanent price increase (there is one seller in their setting). Kaya and Roy (2022, 2024) show that, also under perfect information about the previous transactions, equilibria are also Coasian in the presence of adverse selection, and they find that an upper bound on the buyer's payoff when previous price offers are unobservable (but acceptance decisions are public) is lower than his payoff in some equilibria when price offers are public; they also analyze the effect of increasing competition. Our assumption of imperfect observability of acceptance decisions averts Coasian dynamics, as the buyer's private information is revealed parsimoniously over time in equilibrium, shedding light on the interplay between information and bargaining and providing unique predictions and rich trade dynamics.³ Bonatti and Cisternas (2020) consider a consumer whose willingness to pay evolves stochastically over time and analyze the effects of third-party scores measuring the consumer's purchasing decisions. They find that, in linear Markov equilibria, sellers tend to offer lower prices when scores are less persistent. Our analysis focuses on the case where types are permanent and the information about each transaction is exogenous and persistent (e.g., because it is not transmitted through third parties or because the sellers cannot commit to disregarding some of it). We show that more informative signals tend to be offset by less informative equilibrium purchasing decisions, which may be induced by either low (welfare-enhancing) prices or high (welfare-diminishing) prices.

The work most closely related to ours is Lee and Liu (2013), which studies markets in which traders use bargaining to address disputes, with a type-dependent random outside option being publicly drawn if traders fail to agree. We look at the opposite situation, in which the outside option (value from not trading) is fixed, and learning occurs only through the buyer's endogenous acceptance decisions (so it depends on the buyer's type only through his equilibrium behavior).⁴ We believe this is a more suitable

³ Villas-Boas (2004) studies the case where the monopolist faces overlapping generations of buyers who each live for two periods, thus avoiding Coasian dynamics. He shows that equilibrium involves cycles in the prices offered to new consumers.

⁴ Relatedly, Strulovici (2017) studies a model where a buyer and a seller bargain over quantities of goods, so the buyer has the incentive to conceal she has a high type as it may cause higher prices for subsequent units. He obtains that while Coasian dynamics and efficiency hold for linear production costs, the seller captures part of the surplus of the high type.

framework for studying how information affects prices and trade probability. The implied dynamics and equilibrium mechanisms are therefore significantly different from those of Lee and Liu (2013). In their model, the predictions are similar to those for the analogous model with one trade (Daley and Green, 2012): for intermediate beliefs, all offers are rejected, and for extreme beliefs, either all offers are accepted by both types, or all offers are accepted by exactly one type. In our model, there is either total or partial separation for intermediate beliefs, and the speed of learning is affected not only by the buyer's behavior but also by the endogenous seller's pricing decision. This outcome differs significantly from the equilibrium outcome of the analogous model with one trade.⁵

Our work is also related to the reputation literature involving a privately informed long-lived player interacting with a sequence of short-lived players, initiated by Fudenberg and Levine (1989). Here, the paper closest to ours is that of Faingold and Sannikov (2011), in which a firm sells goods at a fixed price to a continuum of buyers, and information about the firm's previous quality choices is revealed through a diffusion process. The firm either exogenously produces only high-quality goods or can produce either low-quality goods or, at a greater cost, high-quality goods. The price is exogenously set to be equal to 1. Our model endogenizes price formation through a simple (yet canonical) pricing procedure often used in decentralized markets: the seller posts a price, which the buyer accepts or rejects. The model thus incorporates both the buyer's incentives (to reveal/conceal private information by accepting/rejecting high prices) and the sellers' incentives (to choose prices that maximize the expected revenue given the buyer's incentives); these jointly determine the speed of learning. Hence, unlike Faingold and Sannikov (2011), we obtain that the unique equilibrium features (i) efficient discounts and (ii) inefficient rejections. Furthermore, in Faingold and Sannikov (2011), the observability of the individual choices of short-lived agents is irrelevant, whereas in our model the buyer prefers prices to be observable rather than hidden.⁶

The rest of the paper is organized as follows. We present the model with observable prices in Section 2, and we analyze it in Section 3. In Section 4, we study the welfare effects of reducing the information observed by sellers, both by lowering the signal informativeness and by making previous price offers unobservable. In Section 5, we discuss policy implications for several applications and conclude. Appendix A contains the proofs of the results. Appendix B provides results for discrete-time versions of our model.

⁵ It is not difficult to see that the Coase conjecture holds in the one-trade version of our model.

⁶ Note that our stage game is an ultimatum game, which is close to the chain store game (studied by Kreps and Wilson, 1982, and Milgrom and Roberts, 1982). Note also that imperfect information about previous transactions has also been studied in Liu (2011), Liu and Skrzypacz (2014), and Pei (2024b,a), who study reputation effects in repetitions of the product-choice game with limited records. As we argue above, the sequential nature of the ultimatum game—in contrast to the simultaneous nature of the product-choice game, also used by Faingold and Sannikov (2011)—permits studying how endogenous pricing decisions dynamically affect the release of private information. This is because the offered price influences the buyer's action, and together, they determine the information conveyed by the decision to trade.

2 The model

Time is continuous. There is a buyer. At each instant $t \in \mathbb{R}_+$, the buyer meets a short-lived seller, the “ t -seller”, who offers price p_t . The buyer decides the amount of good he purchases from the t -seller, $a_t \in [0, 1]$, at price $p_t a_t$. The buyer values all sellers’ goods equally. His valuation, denoted θ , also referred to as his type, is private, and it is either ℓ or h with $0 < \ell < h$. A natural interpretation is that the buyer’s type is his willingness to pay, which correlates with his wealth or access to alternative purchasing options. Hence, the buyer’s instantaneous payoff from purchasing a mass of a_t units at price p_t is $a_t(\theta - p_t)$, while the t -seller payoff is $a_t p_t$.⁷ The initial probability that the type is h is $\phi_0 \in (0, 1)$.

There is a public signal about the buyer’s previous purchasing decisions. More concretely, at each instant t , the t -seller observes $(X_{t'})_{t' \in [0, t]}$, with

$$X_t := \mu \int_0^t a_{t'} dt' + B_t, \quad (1)$$

where B_t is a normalized Wiener process and $\mu > 0$ is a parameter capturing the precision of the signal. Throughout Sections 2 and 3, we will consider the case where the t -seller also observes the history of price offers made by previous sellers.

We will use ϕ_t to denote the public belief at time t about the buyer’s type being h (given the signal and price histories). We will focus on Markov strategies. For the buyer with type $\theta \in \{\ell, h\}$, an *acceptance strategy* associates to each belief $\phi \in [0, 1]$ and (on- or off-path) price \hat{p} , a value $\alpha_\theta(\phi, \hat{p}) \in [0, 1]$. An *offer strategy* associates to each belief ϕ a distribution $\tilde{\pi}(\phi) \in \Delta(\mathbb{R})$.

Belief dynamics

We define the following two functions:

$$\begin{aligned} \tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) := & \mu(1 - \phi) \phi (\alpha_h(\phi, \hat{p}) - \alpha_\ell(\phi, \hat{p})) \\ & (\hat{\alpha}(\phi, \hat{p}) - \phi \alpha_h(\phi, \hat{p}) - (1 - \phi) \alpha_\ell(\phi, \hat{p})), \end{aligned} \quad (2)$$

$$\tilde{\sigma}(\phi, \hat{p}; \alpha_\ell, \alpha_h)^2 := \mu(1 - \phi)^2 \phi^2 (\alpha_h(\phi, \hat{p}) - \alpha_\ell(\phi, \hat{p}))^2. \quad (3)$$

⁷ We will often interpret a_t as the probability of accepting an indivisible good offered by seller t , for example saying “the buyer accepts for sure” when $a_t = 1$. The interpretation of a_t as a quantity is more convenient in a continuous-time setting because, as we shall see, the buyer will choose $a_t \in (0, 1)$ in a positive measure of time on the equilibrium path for some parameters (in the case of an indivisible good, that would correspond to the buyer mixing in continuous time his acceptance decision). Appendix B argues that, in discrete-time approximations, the two modeling choices provide equivalent results.

These are interpreted as the instantaneous drift and the diffusion parameters of the belief process, respectively, when the sellers assign a probability ϕ to the valuation of the buyer being h , they believe that the ℓ -buyer's acceptance strategy is α_ℓ and the h -buyer's acceptance strategy is α_h , the buyer actual acceptance strategy is $\hat{\alpha}$, and the price offered is equal to \hat{p} .

Equilibrium concept

Definition 1. A (regular Markov perfect) equilibrium is a strategy profile $(\alpha_\ell, \alpha_h, \tilde{\pi})$ together with a pair of value functions $(V_\ell, V_h): [0, 1] \rightarrow \mathbb{R}$ satisfying the following conditions:

- (i) V_ℓ and V_h are continuous, piecewise twice differentiable, and differentiable at all ϕ such that $\alpha_\ell(\phi, \hat{p}) \neq \alpha_h(\phi, \hat{p})$ for some \hat{p} ,⁸ and, for almost all $\phi \in (0, 1)$, they satisfy

$$r V_\theta(\phi) = \mathbb{E}_{\tilde{p} \sim \tilde{\pi}(\phi)} \left[r \alpha_\theta(\phi, \tilde{p})(\theta - \tilde{p}) + \tilde{\mu}(\phi, \tilde{p}; \alpha_\theta, \alpha_\ell, \alpha_h) V'_\theta(\phi) + \frac{1}{2} \tilde{\sigma}(\phi, \tilde{p}; \alpha_\ell, \alpha_h)^2 V''_\theta(\phi) \right], \quad (4)$$

where $\mathbb{E}_{\tilde{p} \sim \tilde{\pi}(\phi)}[\cdot]$ is the expectation operator with respect to the variable \tilde{p} distributed according to $\tilde{\pi}(\phi)$.

- (ii) For all ϕ and \hat{p} , $\alpha_\theta(\phi, \hat{p})$ belongs to

$$\arg \max_{\hat{\alpha} \in [0, 1]} \left(r \hat{\alpha}(\theta - \hat{p}) + \tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) V'_\theta(\phi) \right). \quad (5)$$

- (iii) For all ϕ , $\tilde{\pi}(\phi)$ belongs to

$$\arg \max_{\tilde{\pi} \in \Delta(\mathbb{R})} \mathbb{E}_{\tilde{p} \sim \tilde{\pi}} \left[\left((1 - \phi) \alpha_\ell(\phi, \tilde{p}) + \phi \alpha_h(\phi, \tilde{p}) \right) \tilde{p} \right]. \quad (6)$$

The first condition in Definition 1 requires that the strategy profile and continuation values are regular enough that the Bellman equation (4) holds for both $\theta \in \{\ell, h\}$. The second condition in Definition 1 says that the buyer acts optimally. Equation (5) indicates the tradeoff he faces. If a price \hat{p} is such that $\alpha_\ell(\phi, \hat{p}) = \alpha_h(\phi, \hat{p})$ (i.e., if the signal is uninformative and so $\tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) = 0$), then the buyer accepts for sure any price below his valuation (i.e., any $p < \theta$). If, instead, $\hat{p} \in (\ell, h)$ and $\alpha_\ell(\phi, \hat{p}) < \alpha_h(\phi, \hat{p})$, then acceptance gives the h -buyer an instantaneous payoff equal to $h - \hat{p} > 0$, but reveals information about his type to future sellers, which affects the continuation payoff (as the drift of the posterior is positive).

⁸ Note that if $\alpha_\ell(\phi, \hat{p}) = \alpha_h(\phi, \hat{p})$ for all \hat{p} , then, for any continuation play, the belief remains equal to ϕ . Conversely, if $\alpha_\ell(\phi, \hat{p}) \neq \alpha_h(\phi, \hat{p})$ for some \hat{p} , then the differentiability of the continuation values allows us to compute the buyer's incentive to accept the price offer using equation (5). Note also that if $\alpha_\ell(\phi, \hat{p}) = \alpha_h(\phi, \hat{p})$ for some \hat{p} , then $\tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) = 0$ for all $\hat{\alpha}$; hence the second term in the argument of $\arg \max$ in equation (5) is 0.

The third condition in Definition 1 says that sellers behave optimally. When the posterior is ϕ , the seller chooses the price to maximize the expected revenue, that is, the price multiplied by the probability that it is accepted.

We now present a condition on the equilibrium behavior.

Condition 1. The buyer accepts for sure any offer less than or equal to ℓ .

From now on, we focus on equilibria satisfying Condition 1, which we call simply equilibria. Condition 1 is intuitive, and it is a convenient way to make the analysis tractable. It is analogous to a result obtained in most bargaining models with one purchase: the *Diamond paradox* establishes that, if a buyer receives one offer at a time, the lowest equilibrium offer is no lower than the lowest buyer valuation.^{9,10}

Note that Condition 1 effectively transforms the ℓ -buyer into a “behavior” or “action” type who accepts an offer if and only if it is weakly lower than ℓ . It is then suboptimal for sellers to offer prices below ℓ in equilibrium. As we will see, equilibria under Condition 1 have the property that it is optimal for both types of buyer to behave as prescribed by the condition. Hence, Condition 1 can be seen as either an equilibrium refinement or a behavioral assumption.

3 Equilibrium analysis

3.1 Preliminary results

We begin by presenting some preliminary results that will help build intuition for our main results. Throughout, we will fix an equilibrium $(\alpha_\ell, \alpha_h, \tilde{\pi})$. These will establish some necessary conditions that strategy profiles have to satisfy to be an equilibrium.

Offered prices: We first note that, in the one-shot game, a seller offers ℓ if $\phi < \phi^*$ and h if $\phi > \phi^*$, where $\phi^* := \ell/h$. The same occurs when the signal is uninformative because, when $\mu = 0$, we have $\tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) = 0$, so the buyer behaves myopically (from equation (5)). As the following result establishes, the threshold ϕ^* also plays an important role when the signal is informative.

Lemma 1. *The following hold in any equilibrium:*

(i) *For each $\phi \leq \phi^*$, the support of $\tilde{\pi}(\phi)$ is $\{\ell\}$.*

(ii) *For each $\phi \in (\phi^*, 1)$, the support of $\tilde{\pi}(\phi)$ is either $\{p(\phi)\}$ or $\{\ell, p(\phi)\}$ some $p(\phi) \in (\ell, h)$.*

⁹ In a repeated-trade setting without noise, prices lower than ℓ can be sustained in equilibrium, for example by “punishing” the buyer with prices equal to h if he accepts a higher price. In our model, the noise in the acceptance signal rules out this possibility.

¹⁰ Other authors make other assumptions with similar effects on the equilibrium play. For example, Lee and Liu (2013) require the value functions to be monotone, and the reputation literature assumes that all types except one are behavioral.

Lemma 1 implies that price offers are always smaller than h . This result would be trivial in a model with a unique transaction, since no buyer type would accept an offer larger than h ; hence each seller would be strictly better off offering ℓ than offering a price higher than h . In our model with repeated trade, the result is not obvious, because the h -buyer has signaling motives. For example, it could be that the continuation value of the buyer is increasing in the posterior; hence, by accepting a price offer higher than h , the buyer increases his equilibrium continuation payoff. The proof of Lemma 1 shows that this is not possible: the continuation value of the h -buyer is a decreasing function of the posterior in all equilibria. As a result, the h -buyer never accepts an offer higher than h , as doing so both gives him a negative payoff and decreases his continuation value on expectation.

The previous observation implies that, as in the static model, ℓ is offered and accepted for sure in equilibrium when $\phi < \phi^*$. Indeed, an immediate implication of Condition 1 is that no seller offers a price strictly below ℓ , and hence the equilibrium payoff of the ℓ -buyer is 0. This implies that the ℓ -buyer rejects all offers above ℓ . As a result, when $\phi < \phi^*$, offering ℓ (which is accepted by both types of buyer) gives the seller a larger payoff than offering any price in $(\ell, h]$ (which is rejected by the ℓ -buyer).

The fact that, when $\phi > \phi^*$, the support of the price distribution is either $\{p(\phi)\}$ or $\{\ell, p(\phi)\}$ for some $p(\phi) \in (\ell, h)$ is obtained as follows. Assume a seller offers $\hat{p} \in (\ell, h)$ (on or off path). The h -buyer cannot be strictly willing to reject such a price: if he were, the signal would be deemed uninformative by future sellers (i.e., $\tilde{\mu}(\phi, \hat{p}; \hat{\alpha}, \alpha_\ell, \alpha_h) = 0$), but then the buyer would have the incentive to accept the price (because $\alpha_h(\phi, \hat{p}) = 1$ would be the unique element in the argument of the maximum in (5)). Hence, either the h -buyer is indifferent between accepting the price \hat{p} or not, or he has a strict incentive to accept it. From equation (5), he is indifferent if and only if

$$\underbrace{h - \hat{p}}_{\text{surplus from trade}} = \underbrace{\mu/r(1-\phi)\phi\alpha_h(\phi, \hat{p})(-V'(\phi))}_{\text{reputation loss}} \quad (7)$$

(here and from now on, to save notation, we denote the h -buyer's continuation value by V instead of V_h). The term on the left-hand side of equation (7) is the buyer's instantaneous gain from accepting the offer; it equals his valuation minus the price. The term on the right-hand side is the implied loss in terms of continuation value, which can be interpreted as a reputation loss. Keeping all else equal, this term is larger when the signal is more informative, the equilibrium acceptance probability is larger, or the continuation value is more sensitive to changes in the posterior.

Acceptance decisions: The h -buyer is strictly willing to accept \hat{p} when the left-hand side of equation (7)

is strictly bigger than the right-hand side. This implies that, in equilibrium,

$$\alpha_h(\phi, \hat{p}) = \min \left\{ \frac{r(h-\hat{p})}{\mu(1-\phi)\phi(-V'(\phi))}, 1 \right\}. \quad (8)$$

This acceptance probability acts as a *downward-sloping demand*: it is 1 if \hat{p} is low enough and decreases linearly as \hat{p} increases until it reaches 0, when $\hat{p}=h$. We can then compute the price $\hat{p} \in (\ell, h)$ that maximizes the seller's payoff $\hat{p} \alpha_h(\phi, \hat{p})$. The price $p(\phi)$ in Lemma 1 is the unique maximizer of $\hat{p} \alpha_h(\phi, \hat{p})$, which is given by

$$p(\phi) := \max \left\{ h/2, \overbrace{h - \mu/r(1-\phi)\phi(-V'(\phi))}^{(*)} \right\}. \quad (9)$$

The expression $(*)$ in equation (9) represents the highest price accepted with probability one by the h -buyer, which corresponds to the kink of $\alpha_h(\phi, \hat{p})$. The seller's optimal offer is then either the maximizer of the linear part of $\alpha_h(\phi, \hat{p})$ —that is, $h/2$, which is rejected by the h -buyer with positive probability—or the corner solution $(*)$, if that price is above $h/2$ —which is accepted by the h -buyer for sure. In particular, we have

$$\alpha(\phi) < 1 \Rightarrow p(\phi) = h/2, \quad (10)$$

where from now on $\alpha(\phi) := \alpha_h(\phi, p(\phi))$ is the equilibrium probability that the h -buyer accepts $p(\phi)$ (recall that such an offer is rejected for sure by the ℓ -buyer).

It is important to note that, in any equilibrium, it is optimal for the h -buyer to mimic the ℓ -buyer, that is, to accept an offer if and only if it is equal to or lower than ℓ . This follows from the observation that either the high price $p(\phi)$ is the largest price that the buyer accepts with probability one, in which case he is indifferent between acceptance and rejection, or $p(\phi)$ equals $h/2$, and the buyer randomizes between acceptance and rejection. Thus, (7) holds for $\hat{p} = p(\phi)$ and $\alpha_h(\phi, \hat{p}) = \alpha(\phi)$.

Optimal high prices: The gain a seller obtains from offering $p(\phi)$ is $\phi \alpha(\phi) p(\phi)$. By Lemma 1, it is weakly optimal for a seller to offer $p(\phi) > \ell$ for all $\phi > \phi^*$. Also, offering ℓ gives a seller a payoff equal to ℓ (since she sells for sure). Hence, we have that

$$\phi \alpha(\phi) p(\phi) \geq \ell \quad (11)$$

for all $\phi > \phi^*$. The previous expression holds with equality when ℓ is offered with positive probability. Since, by equation (10), either $\alpha(\phi) = 1$ or $p(\phi) = h/2$ (or both), we have the following condition for

sellers to offer $p(\phi)$ with positive probability for $\phi > \phi^*$:

$$\pi(\phi) \in (0, 1) \Rightarrow p(\phi) = \max\{h/2, \ell/\phi\}, \quad (12)$$

where, by another abuse of notation, $\pi(\phi)$ indicates the probability that the offered price equals $p(\phi)$ (in this case, the term $(*)$ in equation (9) is equal to ℓ/ϕ). Hence, ℓ is offered with probability $1 - \pi(\phi)$ (by Lemma 1).

3.2 Equilibrium characterization

Less informative signal

We first focus on the case where the signal is relatively uninformative—that is, the case where μ is small (in a sense that will be made precise). Equivalent results can be obtained when the buyer is relatively impatient, that is, in the case where r is large (note that the equilibrium objects, $\{V, p, \pi, \alpha\}$, determined through equations (4)-(6), depend on r and μ only through r/μ). Note that, from the previous section, a strategy profile is fully determined by $p(\phi)$ (the high price intended only for the h -buyer), $\pi(\phi)$ (the probability with which a seller offers $p(\phi)$), and $\alpha(\phi)$ (the probability with which the h -buyer accepts $p(\phi)$), for each $\phi \in [0, 1]$.

The following proposition states some important properties of the unique equilibrium. Below, we provide intuition for why these properties must hold, and we argue that a discount mechanism is necessary to balance the buyer's and sellers' incentives.

Proposition 1. *There is some largest $\bar{\mu} \in (0, +\infty]$ such that, for each $\mu < \bar{\mu}$, there is an essentially unique equilibrium.¹¹ In every such equilibrium, there is some $\phi^\dagger \in (\phi^*, 2\phi^*]$ such that the following hold:*

- (i) *On $(0, \phi^*]$, π is equal to 0.*
- (ii) *On (ϕ^*, ϕ^\dagger) , π is strictly increasing, α is equal to 1, and p is strictly decreasing.*
- (iii) *On $(\phi^\dagger, 1)$, π and α are equal to 1, and p is strictly increasing.*

Figure 1 illustrates Proposition 1. For $\phi < \phi^*$, the sellers offer ℓ , the signal is uninformative in equilibrium, and the payoff of the h -buyer is $h - \ell$. As in the static game, the sellers are pessimistic

¹¹ “Essentially unique” here means that other equilibria differ from it only on a zero-measure set of posteriors that do not affect the outcome of the game.

¹² Note that ℓ is offered with probability one for all $\phi \in (0, \phi^*)$ (i.e., $\alpha(\phi) = 1$). Hence, even though p is not defined in this region (recall Lemma 1), we think it is natural to depict it as $p(\phi) = h$ for all $\phi \in (0, \phi^*)$, as it is the price that gives highest payoff to the sellers conditional on offering a price above ℓ .




Figure 1: Various equilibrium objects for $h=r=1$, $\ell=0.3$, and $\mu=0.8$.¹²

enough about the buyer's valuation that they offer price ℓ even though the h -buyer is willing to accept any price below h .

Now consider a posterior ϕ higher than, but close to, the threshold ϕ^* . Note that the seller's payoff from offering a price $\hat{p} \in (\ell, h)$ is $\phi \alpha_h(\phi, \hat{p}) \hat{p}$, which is higher than ℓ only if $\hat{p} \geq \ell/\phi$. Hence, $p(\phi) \geq \ell/\phi$, so $p(\phi)$ is close to h if ϕ is close to ϕ^* . By (10), $p(\phi)$ is accepted for sure by the h -buyer. It may seem contradictory that, when the posterior is close to ϕ^* , the buyer is willing to accept high prices—which give him little surplus—even though rejection would be highly informative and would bring the posterior close to ϕ^* , where his continuation value is maximal. The apparent contradiction is explained by a *discount mechanism*: in equilibrium, sellers offer ℓ with a high probability (i.e., $\pi(\phi) \in (0, 1)$), which implies that necessarily $p(\phi) = \ell/\phi$ (by (12)) for ϕ close to ϕ^* . As we see in Figure 1, such offers flatten the buyer's continuation value for posteriors close to ϕ^* ; consequently, the signaling gain from rejection is small enough that the h -buyer is willing to accept high prices. In particular, because the h -buyer's payoff from accepting ℓ/ϕ vanishes as ϕ gets close to ϕ^* (since ℓ/ϕ gets close to h), $V'(\phi)$ tends to 0 as $\phi \rightarrow \phi^*$ (recall equation (7)), so V is smooth at ϕ^* .

We now argue that, for high posteriors, a high price $p(\phi) > h/2$ is offered for sure, and it is accepted for sure by the h -buyer. To see this, we change variables and use the log-likelihood $\check{z}(\phi) := \log(\phi/(1-\phi))$; we set $W(\check{z}(\phi)) := V(\phi)$.¹³ Equations (8) and (9) can then be written as

$$\alpha_h(\phi, \hat{p}) = \min \left\{ \frac{r(h-\hat{p})}{\mu(-W'(\check{z}(\phi)))}, 1 \right\} \quad \text{and} \quad p(\phi) = \max \left\{ h/2, h - \mu/r(-W'(\check{z}(\phi))) \right\}, \quad (13)$$

respectively. Since W is decreasing and bounded and $\lim_{\phi \rightarrow 1} \check{z}(\phi) = +\infty$, it must be that $\lim_{\phi \rightarrow 1} W'(\check{z}(\phi)) =$

¹³ The log-likelihood is often more convenient to use than the posterior. For example, when the type-dependent acceptance probabilities are constant, the drift of the log-likelihood is constant, while the drift of the posterior is proportional to $\phi(1-\phi)$ because learning is slow for extreme posteriors.

0. Hence, $p(\phi) \rightarrow h$ as $\phi \rightarrow 1$, and so $\alpha(\phi) = 1$ for large enough ϕ (by (10)). Also, from (12) we have that $\pi(\phi) = 1$ for ϕ large enough.

The proof of Proposition 1 establishes that when μ is small, only equilibria with the two regimes described above can exist.¹⁴ That is, there is a threshold $\phi^\dagger \in (\phi^*, 1)$ at which sellers stop offering discounts. Such a threshold is obtained using the standard smooth-pasting condition at ϕ^\dagger so that the differentiability of V is preserved.

Since $p(\phi) = \ell/\phi > h/2$ for $\phi < \phi^\dagger$, we have

$$\ell/\phi^\dagger = p(\phi^\dagger) \geq h/2 \Rightarrow \phi^\dagger \leq 2\phi^* .$$

In fact, ϕ^\dagger is increasing in μ (see the proof of Proposition 3) and reaches $2\phi^*$ when $\mu = \bar{\mu}$. As a result, in the case $\mu > \bar{\mu}$ (analyzed below), there is no equilibrium like the one described in Proposition 1. Note that $\bar{\mu} = +\infty$ if and only if $\phi^* \geq 1/2$ (i.e., $h/2 \leq \ell$).

Remark 3.1. We can see in Figure 1 that $\pi(\cdot)$ is discontinuous at ϕ^\dagger . To see why, note first that when $\phi \in (\phi^*, \phi^\dagger)$, the price is $p(\phi) = \ell/\phi$, which is decreasing. Hence, for the buyer to remain indifferent in equilibrium between accepting and rejecting $p(\phi)$, his continuation value has to decrease faster for higher values of ϕ (see equation (13)). In fact, in the proof of Proposition 1 we show that $W''(\check{z}(\phi)) = -(r/\mu)\ell e^{-\check{z}(\phi)} < 0$ for $\phi \in (\phi^*, \phi^\dagger)$, that is, W is strictly concave in this region. Because only high prices are offered in the region $(\phi^\dagger, 1)$, which are accepted for sure by the h -buyer, the log-likelihood follows a standard diffusion process with constant drift. Hence, using that rejecting all offers above ℓ is optimal for the h -buyer, his continuation value can be shown to be $W(\check{z}(\phi)) = c_1 e^{-\kappa\check{z}(\phi)} > 0$ for some constants $c_1, \kappa > 0$, which is strictly convex. At ϕ^\dagger , equation (4) implies

$$\begin{aligned} rW(\phi^\dagger) &= \lim_{\phi \uparrow \phi^\dagger} (1 - \pi(\phi))r(h - \ell) + \pi(\phi) \frac{\mu}{2} (-W'(\check{z}(\phi)) + W''(\check{z}(\phi))) \\ &= \lim_{\phi \downarrow \phi^\dagger} \frac{\mu}{2} (-W'(\check{z}(\phi)) + W''(\check{z}(\phi))) . \end{aligned}$$

Hence, because W is differentiable at ϕ^\dagger and the right and left limits of $W''(\check{z}(\phi))$ as ϕ approaches ϕ^\dagger differ, we have that $\lim_{\phi \uparrow \phi^\dagger} \pi(\phi) < 1 = \lim_{\phi \downarrow \phi^\dagger} \pi(\phi)$.

¹⁴ This is intuitive: from (10) we have that $\alpha(\phi) < 1$ only if $p(\phi) = h/2$, but such a low price is not reached in equilibrium when μ is small. In fact, as $\mu \rightarrow 0$, we have that $p(\phi) \rightarrow h$ for all $\phi > \phi^*$; that is, the equilibrium outcome converges to the static equilibrium outcome as the signal becomes less informative.




Figure 2: Various equilibrium objects for $h=r=1$, $\ell=0.15$, and $\mu=1.5$.

More informative signal

We now study the case where the signal is relatively informative, that is, where $\mu > \bar{\mu}$. This is equivalent to studying the case where the buyer is relatively patient. As discussed above, this case occurs only if $\phi^* < 1/2$ (since $\bar{\mu} = +\infty$ otherwise).

As we argued above, there are no equilibria with the structure provided in Proposition 1 when $\mu > \bar{\mu}$. The following result establishes that, when the signal is informative, there is again an essentially unique equilibrium, which now features two new equilibrium regions in addition to the three regions described in Proposition 1. These regions are depicted in Figure 2. Section 4.1 characterizes the limit $\mu \rightarrow \infty$

Proposition 2. *Let $\bar{\mu}$ be as defined in Proposition 1. Then, for all $\mu > \bar{\mu}$, there is an essentially unique equilibrium. For each $\mu > \bar{\mu}$ there are thresholds $\hat{\phi}^\dagger \in (2\phi^*, 1)$ and $\hat{\phi}^{\dagger\dagger} \in (\hat{\phi}^\dagger, 1)$ such that, in the unique equilibrium, the following hold:*

- (i) On $(0, \phi^*)$, π is equal to 0.
- (ii) On $(\phi^*, 2\phi^*)$, π is strictly increasing, α is equal to 1, and p is strictly decreasing.
- (iii) On $(2\phi^*, \hat{\phi}^\dagger)$, π is strictly increasing, α is strictly decreasing, and p is equal to $h/2$.
- (iv) On $(\hat{\phi}^\dagger, \hat{\phi}^{\dagger\dagger})$, π is equal to 1, α is strictly increasing, and p is equal to $h/2$.
- (v) On $(\hat{\phi}^{\dagger\dagger}, 1)$, π and α are equal to 1, and p is strictly increasing.

As before, there is a lower region $(0, \phi^*)$ where ℓ is offered and accepted for sure. For low-intermediate posteriors (posteriors close to but above ϕ^*), sellers randomize between offering ℓ , which both types of buyer accept for sure, and offering $p(\phi) = \ell/\phi$, which the h -buyer accepts for sure (recall (12)). As the posterior increases, the high price $p(\phi)$ decreases. Now, since $\mu > \bar{\mu}$, there is no $\phi^\dagger \in (\phi^*, 1)$ where the equilibrium transition to a phase where $\alpha(\phi) = \pi(\phi) = 1$: as we argued before, if such a value ϕ^\dagger existed, it would be larger than $2\phi^*$, but this would imply that $p(\phi^\dagger) < h/2$, contradicting (9).

Hence, when the posterior reaches $2\phi^*$, the price is equal to $h/2$, which is the lowest price above ℓ offered in equilibrium (by (9)). Then, for posteriors close to but higher than $2\phi^*$, additional increases in the posterior do not decrease the price $p(\phi)$ further: it remains equal to $h/2$. As ϕ increases, the sellers' indifference between offering ℓ and offering $h/2$ is maintained by a *rejection mechanism*: an increase in the posterior lowers the probability with which $h/2$ is accepted. The sellers' indifference condition requires that $\phi \alpha(\phi)h/2 = \ell$ or, equivalently, $\alpha(\phi) = 2\ell/(\phi h)$, so α decreases in ϕ . As the posterior increases, the probability with which sellers offer ℓ decreases. Both the positive probability of a discount and the positive probability of rejection of $h/2$ slow the sellers' learning, keeping the h -buyer indifferent between accepting and rejecting $h/2$.¹⁵

As for the case where $\mu < \bar{\mu}$, sellers stop offering ℓ when the posterior reaches a certain threshold, denoted by $\hat{\phi}^\dagger > 2\phi^*$. While π jumps at $\hat{\phi}^\dagger$ for the same reason as before, α is continuous at $\hat{\phi}^\dagger$. Intuitively, because p is continuous (by (9)), a continuous α ensures that the sellers continue to find it optimal to offer $p(\phi)$. By (10), the price remains equal to $h/2$ in some range of higher posteriors, denoted by $(\hat{\phi}^\dagger, \hat{\phi}^{\dagger\dagger})$, where α is increasing in ϕ . Intuitively, as ϕ increases and the posterior region with discounts gets further away, so the signaling incentive for the h -buyer to reject $h/2$ decreases. Such incentive is kept in equilibrium by increasing the informativeness of rejection. After α reaches 1, it stays equal to 1. As we argued in Section 3.1, the h -buyer stays indifferent between accepting and rejecting the high price p ; hence, the price increases in the posterior.

4 Effects of privacy policies

We now study the effects of reducing the amount of information each seller has about the previous history. The results give insight into the possible impact of policies regulating or banning cookies, which we discuss further in Section 5.

In our model, there are two ways of reducing the information available to each seller. We will first study the effect of making the acceptance signal less precise (i.e., of reducing μ) in Section 4.1. We will then analyze a model in which sellers can observe the acceptance signal but not the prices offered by previous sellers in Section 4.2. We will finally study the welfare effect of making the price offers unobservable in Section 4.3.

¹⁵ The proof of Proposition 1 shows that ϕ^\dagger increases as μ increases when the signal is less informative—indicating that it becomes easier for the buyer to lower the seller's posterior—and $\bar{\mu}$ is such that $\phi^\dagger = 2\phi^*$. It then follows that the rejection mechanism only can occur when the signal is informative enough, that is, when $\mu > \bar{\mu}$.

4.1 Limiting signal precision

We first study the effect of reducing signal informativeness. Reductions in μ could correspond to policies limiting the data stored in cookies, while increases in μ could be attributed to improvements in tracking technology or to regulations requiring transparency in the transactions of government agencies.

Proposition 3. (i) For each $\phi \in (\phi^*, 1)$, $V(\phi)$ is strictly increasing in μ on $(0, \bar{\mu})$.

(ii) If $\mu > \bar{\mu}$ and $\phi \in (\phi^*, \hat{\phi}^\dagger)$, then $\frac{d}{d\mu}V(\phi) > 0$.

(iii) For each $\phi \in [0, 1)$, $\lim_{\mu \rightarrow \infty} V(\phi) = h - \ell$.

(iv) For each $\phi \in [0, 1]$, $\lim_{\mu \rightarrow 0} V(\phi) = (h - \ell)\mathbb{I}_{[0, \phi^*]}(\phi)$.

The first claim in Proposition 3 establishes that when the signal is not very informative, the h -buyer's payoff is increasing in the signal precision. The intuition for the result is the following. When the signal becomes more informative, the buyer's acceptance of a high price is more informative about his high valuation. Greater signal informativeness is offset, in equilibrium, by more frequent discounts, which increase—and therefore flatten—the h -buyer's continuation value. When the signal informativeness is low, only this discount effect takes place, thus increasing the signal informativeness benefits the buyer.

When the signal is already informative, it is less clear whether informativeness translates into a higher buyer payoff or not. Proposition 3 establishes that small increases in μ increase $V(\phi)$ when ϕ is in $(\phi^*, \hat{\phi}^\dagger]$, so the strengthening of the discount mechanism dominates the strengthening of the rejection mechanism in this region. Since only the rejection mechanism takes place in $(\hat{\phi}^\dagger, \hat{\phi}^{\dagger\dagger})$, higher signal precision makes it more costly for the h -buyer to mimic the ℓ -buyer and reach $\hat{\phi}^\dagger$ (where discounts begin). Still, since the continuation value upon reaching $\hat{\phi}^\dagger$ increases, the change in $V(\phi)$ for $\phi \in (\hat{\phi}^\dagger, 1)$ is unclear. While the proof of Proposition 3 provides explicit equations for $\hat{\phi}^\dagger$, $\hat{\phi}^{\dagger\dagger}$, and the continuation value for the different regions, analytically determining the effect of changes in μ has proved not possible. Numerical simulations seem to indicate that increases in μ do increase $V(\phi)$ for all $\phi \in (\phi^*, 1)$.

When the signal is very imprecise (small μ), the h -buyer obtains a very low payoff for most $\phi > \phi^*$: the region where sellers offer discounts vanishes as $\mu \rightarrow 0$ and, due to slow learning, mimicking the ℓ -buyer becomes unattractive. By contrast, when the signal is very precise (large μ), the h -buyer's obtains a payoff close to $h - \ell$ for most $\phi > \phi^*$. As μ increases, the region where sellers offer discounts expands toward $(\phi^*, 1)$, and the probability of a discount tends to 1 (i.e., $\pi(\phi) \rightarrow 0$ as $\mu \rightarrow \infty$ for all $\phi \in (\phi^*, 1)$).

4.2 The private offers case

We now study the case where offers are unobservable to future sellers, which is referred to in the literature as the “private offers case” (the model presented and studied in Sections 2 and 3 corresponds to the “public

offers case”, see Hörner and Vieille, 2009).

In practice, price offers may be unobservable because of regulations restricting the information available to sellers. For example, cookies may be allowed to collect metrics about a user’s previous activity but not data on his actual transactions. A search engine may be able to track a user up to the point where he opens a webpage and collect some information about his behavior on the webpage, but not to identify the actual offers made to him. Similarly, while the existence of previous procurement contracts may be easier to be known by current contractors, some details on prices may be undisclosed under confidentiality agreements.

We now construct a version of the model described in Section 2 in which price offers are unobservable to other sellers. Now, the t -seller observes only the public signal $(X_{t'})_{t' < t}$ defined in (1). Markov strategies are defined in exactly the same way as in Section 2. Given a strategy profile, for each $\theta \in \{\ell, h\}$ we define the θ -buyer’s expected acceptance probability for the belief ϕ as

$$\bar{\alpha}_\theta(\phi) = \mathbb{E}_{\tilde{p} \sim \tilde{\pi}(\phi)}[\alpha_\theta(\phi, \tilde{p})].$$

The value of $\bar{\alpha}_\theta(\phi)$ indicates the equilibrium probability with which a θ -buyer accepts the price offer when the posterior is ϕ . Hence, the Bellman equation is now

$$\begin{aligned} r V_\theta(\phi) = & \mathbb{E}_{\tilde{p} \sim \tilde{\pi}(\phi)}[r \alpha_\theta(\phi, \tilde{p})(\theta - \tilde{p}) + \tilde{\mu}(\phi, \tilde{p}; \alpha_\theta, \bar{\alpha}_\ell, \bar{\alpha}_h) V'_\theta(\phi) \\ & + \frac{1}{2} \tilde{\sigma}(\phi, \tilde{p}; \bar{\alpha}_\ell, \bar{\alpha}_h)^2 V''_\theta(\phi)] \end{aligned} \quad (14)$$

instead of (4), where $\tilde{\mu}$ and $\tilde{\sigma}$ are defined in (2) and (3). The crucial difference between the analysis of the private offers case and that of the public offers case is that now $\bar{\alpha}_\ell$ and $\bar{\alpha}_h$ do not depend on the actual price offer. Thus, while the drift in equation (14) only depends on \tilde{p} through the acceptance decision of the buyer (i.e., $\alpha_\theta(\phi, \tilde{p})$), while the variance is independent of \tilde{p} .

The equilibrium concept is analogous to that of Definition 1, with the following differences. First, because the price offers are not observable, the condition that V_θ is “differentiable at all ϕ such that $\alpha_\ell(\phi, \hat{p}) \neq \alpha_h(\phi, \hat{p})$ for some \hat{p} ” is now replaced by “ V_θ is differentiable at all ϕ such that $\bar{\alpha}_h(\phi) \neq \bar{\alpha}_\ell(\phi)$.” Second, instead of the condition (5), $\alpha_\theta(\phi, \hat{p})$ belongs to the following set for all ϕ where V_θ is differentiable:

$$\arg \max_{\hat{\alpha} \in [0,1]} (r \hat{\alpha}(\theta - \hat{p}) + \tilde{\mu}(\phi; \hat{\alpha}, \bar{\alpha}_\ell, \bar{\alpha}_h) V'_\theta(\phi)). \quad (15)$$

Equilibrium analysis

We begin by stating that Lemma 1 also holds for the unobservable offers model.

Lemma 2. *Lemma 1 holds in the private offers case.*

As in the public offers case, if $\phi_t \leq \phi^*$ then the t -seller offers ℓ , while if $\phi_t > \phi^*$ she may randomize between offering some price $p(\phi_t)$ (with some probability again denoted by $\pi(\phi_t)$) and offering ℓ (with probability $1 - \pi(\phi_t)$). Again, $\alpha(\phi_t)$ denotes the probability with which the h -buyer accepts the price $p(\phi_t)$ when the t -seller offers it, and V is the continuation value of the h -buyer.

Although, just as in the public offers case, the support of prices consists of either one or two points, the logic for this is quite different in the private offers case. When prices are observable, each price is accepted with a different probability in equilibrium (that is, α_h depends on both \hat{p} and ϕ in equation (7)). This probability affects the informativeness of the signal so that the buyer is indifferent between accepting and rejecting the price (recall equation (7)). As we saw, there is a unique price above ℓ that maximizes the seller's payoff (i.e., the acceptance probability multiplied by the price). When instead prices are unobservable, the buyer's reputation loss from accepting an offer \hat{p} is independent of the price offered. Indeed, the h -buyer is indifferent between accepting and rejecting $\hat{p} \in (\ell, h)$ if

$$\underbrace{h - \hat{p}}_{\text{surplus from trade}} = \overbrace{\mu/r(1-\phi)\phi(\bar{\alpha}_h(\phi) - \bar{\alpha}_\ell(\phi))(-V'(\phi))}_{\text{reputation loss}}, \quad (16)$$

and he is strictly willing to accept (reject) the offer if the right-hand side of (16) is strictly higher (lower) than the left-hand side.

Note that $\bar{\alpha}_\ell(\phi) = 1 - \pi(\phi)$; that is, the ℓ -buyer's acceptance probability coincides with the probability with which ℓ is offered. By the standard take-it-or-leave-it offer argument, a seller offers a price higher than ℓ in equilibrium only if the h -buyer is indifferent between accepting or rejecting it and he accepts it for sure. We then have that $\bar{\alpha}_h(\phi) = 1$. Hence, in equilibrium, sellers offer either ℓ or

$$p(\phi) = h - \mu r^{-1}(1 - \phi)\phi\pi(\phi)(-V'(\phi)). \quad (17)$$

Any off-path offer in $(\ell, p(\phi))$ is accepted for sure by the h -buyer, while any price strictly above $p(\phi)$ is rejected for sure. It then follows that, unlike in the public offers case, the equilibrium probability that the h -buyer accepts $p(\phi)$ (again denoted by $\alpha(\phi)$) is always 1. Hence, the rejection mechanism is not present in the private offers case.

Proposition 4. *Assume price offers are unobservable. Then there is a unique equilibrium. In this equilibrium, there is some $\phi^\ddagger \in (\phi^*, 1)$ such that the following hold:*

(i) *On $(0, \phi^*]$, π is equal to 0.*

(ii) *On (ϕ^*, ϕ^\ddagger) , π is strictly increasing, α is equal to 1, and p is strictly decreasing.*

(iii) *On $(\phi^\ddagger, 1)$, π and α are equal to 1, and p is strictly increasing.*

Proposition 4 resembles Proposition 1, but applies to all values of μ . For the same reasons as in the public offers case, for posteriors in $(0, \phi^*)$, sellers offer ℓ with probability one. Again, when ϕ is close to (but above) ϕ^* , (i) sellers are only willing to offer $p(\phi)$ if $p(\phi) \geq \ell/\phi_t$; and (ii) the h -buyer, in turn, is willing to accept a high price only if his continuation value is not very sensitive to the posterior. The “flattening” of h -buyer’s continuation value is produced by a variation of the previous discount mechanism: Because ℓ is offered with positive probability (i.e., $\pi(\phi) \in (0, 1)$), V stays high and, since offers made by sellers are not observable to future sellers, the signal becomes less informative. Hence, $V'(\phi)$ and $\tilde{\mu}(\phi; \hat{\alpha}, \bar{\alpha}_\ell, \bar{\alpha}_h)$ are small in absolute value, so accepting a high price is more attractive to the h -buyer by (15). Now, π increases in ϕ until it is equal to 1 for some value $\phi = \phi^\ddagger$ (so it is continuous). From then on, as before, only $p(\phi)$ is offered, and $p(\phi)$ increases toward h as $\phi \rightarrow 1$.

We finalize the section with a result analogous to Proposition 3, establishing that a more informative signal is always beneficial for the buyer when prices are not observable.

Proposition 5. *When prices are not observable:*

(i) *For each $\phi \in (\phi^*, 1)$, $V(\phi)$ is strictly increasing in μ .*

(ii) *For each $\phi \in [0, 1)$, $\lim_{\mu \rightarrow \infty} V(\phi) = h - \ell$.*

(iii) *For each $\phi \in (0, 1)$, $\lim_{\mu \rightarrow 0} V(\phi) = (h - \ell) \mathbb{I}_{[0, \phi^*]}(\phi)$.*

4.3 Welfare analysis

In this section, we compare the buyer welfare and efficiency in the public and private offers cases analyzed in Sections 3 and 4.2. We will also consider the benchmark case where the signal is uninformative, interpreted as an online market where cookies are banned and thus buyers are fully anonymous.¹⁶

We denote the three cases as follows: “ $x = \text{no}$ ” refers to the case where the sellers observe neither the prices nor the signal, “ $x = \text{ob}$ ” refers to the case where the sellers observe both the prices and the signal,

¹⁶ There is a movement aimed at banning the practices that allow advertisers and political organizations to track individuals with tailored messages, a practice called “microtargeting”. See, for example, <https://www.politico.eu/article/targeted-advertising-tech-privacy>.

and “ $x = \text{un}$ ” refers to the case where the sellers observe the signal but not the prices. As discussed above, in the case where nothing is observable, the sellers offer ℓ for sure when $\phi < \phi^*$ and offer h for sure when $\phi > \phi^*$.

The buyer’s surplus

We first compare the buyer’s surplus across the different cases. Given that the ℓ -buyer obtains no surplus in any of the cases, we focus on the payoff of the h -buyer.

Proposition 6. (i) If $\mu \leq \bar{\mu}$, then $V^{\text{ob}}(\phi) > V^{\text{un}}(\phi) > V^{\text{no}}(\phi)$ for all $\phi \in (\phi^*, 1)$.

(ii) If $\mu > \bar{\mu}$, then $\min\{V^{\text{ob}}(\phi), V^{\text{un}}(\phi)\} > V^{\text{no}}(\phi)$ for all $\phi \in (\phi^*, 1)$.

Comparing the three cases is easier when the informativeness of the signal is low. The equilibrium structure is then similar in the observable case (Proposition 1) and the unobservable case (Proposition 4). For low posteriors, the sellers offer low prices; for intermediate posteriors, they randomize between low and high prices; and for high posteriors, they offer only high prices. The h -buyer accepts all offers.

For low and high posteriors, the incentives of the buyer and the sellers do not qualitatively depend on the observability of the price offers. Sellers offer the discounted price ℓ when ϕ is low and the highest price which is accepted for sure by the h -buyer when ϕ is high. Let us now fix some intermediate posterior ϕ . In the public offers case, accepting the high price ℓ/ϕ leads to a significant expected increase in the posterior. The buyer nevertheless accepts the high price, because his continuation value is not sensitive to the posterior. On the other hand, in the private offers case, other sellers do not know whether the high price ℓ/ϕ or the low price ℓ was offered. Hence, the buyer’s acceptance of ℓ/ϕ leads to a lower expected increase in the posterior. The continuation value is therefore steeper in the private offers case. Formally, equations (9) and (17) can be written as

$$-\frac{d}{d\phi} V^x(\phi) = \gamma^x(\phi) \frac{r(h-\ell/\phi)}{\mu(1-\phi)\phi}, \quad (18)$$

where $\gamma^{\text{ob}}(\phi) = 1$ and $\gamma^{\text{un}}(\phi) = \frac{1}{\pi^{\text{un}}(\phi)} > 1$. For higher ϕ , V^x solves the same equation for both $x = \text{ob}$ and $x = \text{un}$, since in both cases the sellers make only high offers, and the h -buyer is indifferent between acceptance and rejection. The proof of Proposition 6 shows that this implies that the buyer is indeed better off in the public offers case.

When the signal is informative (i.e., $\mu > \bar{\mu}$), the relative order between V^{ob} and V^{un} may be reversed at some posteriors. The reason lies in the different mechanisms that slow learning in equilibrium. When prices are unobservable, only the discount mechanism is present: the signal becomes less informative only when the low price is offered with positive probability. When prices are observable, learning is slowed




Figure 3: (a) V^{ob} and V^{un} for $h=r=1$, $\ell=0.05$, and $\mu=3$. (b) Thresholds for the public offers case (ϕ^\ddagger , $\hat{\phi}^\ddagger$, and $\hat{\phi}^{\ddagger\ddagger}$) and the private offers case (ϕ^\ddagger) as a function of μ , for $h=r=1$ and $\ell=0.2$.

down by the rejection mechanism as well as the discount mechanism. That is, as described in Section 3.2, the h -buyer sometimes rejects high-price offers, which slows learning for intermediate posteriors. This means that, if the initial prior is high, it takes a long time for the posterior to reach low values when the h -buyer mimics the ℓ -buyer. The implication is that when the signal is informative enough, at high posteriors the buyer is worse off in the public offers case than in the private offers case.

Figure 3(a) depicts V^{ob} and V^{un} for parameters under which, when the signal is very informative, at some posteriors the buyer is worse off in the public offers case than in the private offers case. For intermediate posteriors, sellers offer the low price more often in the public offers case, and so the buyer's payoff is higher there. However, for higher posteriors, sellers offer higher prices in the public offers case. We can see this from the figure: the low acceptance probability at intermediate posteriors implies that V^{ob} decreases quickly to preserve the buyer's indifference between accepting high prices or rejecting them. This means that if the posterior is initially very high, it takes a long time for it to become low, even if the h -buyer mimics the ℓ -buyer. As a result, at high posteriors, sellers offer higher prices in the public offers case, which means the buyer's payoff is lower.

Figure 3(b) depicts the equilibrium thresholds in both the public and private offers cases as functions of the signal precision μ . When the signal is not very informative (i.e., $\mu < \bar{\mu}$), the range where sellers offer discounts in the public offers case, $(0, \phi^\ddagger)$, is larger than its counterpart in the private offers case, $(0, \hat{\phi}^\ddagger)$. While this is also true when the signal is very informative (with discounts offered on $(0, \hat{\phi}^\ddagger)$ in the public offers case and on $(0, \phi^\ddagger)$ in the private offers case), for these signals there is also, in the public offers case, a region $(\hat{\phi}^\ddagger, \hat{\phi}^{\ddagger\ddagger})$ with no discounts and where the buyer rejects high prices in equilibrium, making him worse off than in the private offers case.

Remark 4.1. Proposition 6 establishes that the buyer tends to prefer public offers over private offers in repeated bargaining, as rejecting a high offer then sends a stronger signal of low type. The logic would be reversed if there were only one transaction (see Kaya and Liu, 2015), as in that case unobservability of offers would enhance the Coasian forces, since sellers could not update their beliefs according to previous sellers' deviations. However, our result also shows that unobservability is sometimes beneficial to the buyer in repeated bargaining, since it prevents the rejection mechanism from slowing equilibrium learning.

Efficiency

We now take the perspective of a social planner who has the same discount rate as the buyer. The social planner values each transaction of the θ -buyer at θ , independently of the transaction price, for both $\theta \in \{\ell, h\}$.

We use $W^x(\phi)$ to denote the social welfare (or efficiency) for each case $x \in \{\text{no}, \text{un}, \text{ob}\}$. The social welfare is given by

$$W^x(\phi) = (1-\phi) \overbrace{\mathbb{E}^{x,\ell} \left[\int_0^\infty \ell (1-\pi^x(\phi_t)) e^{-rt} r dt \mid \phi_0 = \phi \right]}^{(*)} + \phi \underbrace{\mathbb{E}^{x,h} \left[\int_0^\infty h (1-\pi^x(\phi_t) + \pi^x(\phi_t) \alpha^x(\phi_t)) e^{-rt} r dt \mid \phi_0 = \phi \right]}_{(**)} \quad (19)$$

for all $\phi \in (\phi^*, 1)$, where $\mathbb{E}^{x,\theta}$ is the expectation in the equilibrium for the x -model conditional on the strategy of the θ -buyer.¹⁷ The term $(*)$ is equal to the social welfare generated by the transactions of the ℓ -buyer, who purchases only when the price is ℓ . The term $(**)$ is equal to the social welfare generated by the transactions of the h -buyer, who purchases for sure when the price is ℓ and with probability $\alpha^x(\phi)$ otherwise.

Proposition 7. (i) If $\mu \leq \bar{\mu}$, then $W^{\text{ob}}(\phi) > W^{\text{un}}(\phi) > W^{\text{no}}(\phi)$ for all $\phi \in (\phi^*, 1)$.

(ii) If $\mu > \bar{\mu}$, then $W^{\text{un}}(\phi) > W^{\text{no}}(\phi)$ for all $\phi \in (\phi^*, 1)$.

We observe that the term $(*)$ in equation (19) is 0 when $x = \text{no}$ and $\phi > \phi^*$. Hence, the social welfare gain from the presence of a signal (with or without information about prices) comes from the fact that sellers then offer ℓ more frequently than if there is no signal, which implies that the ℓ -buyer purchases more often. Recall that the h -buyer's payoff from mimicking the ℓ -buyer coincides with his equilibrium payoff, since he is always indifferent between accepting and rejecting $p(\phi)$. Note also that, if the h -buyer

¹⁷ Note that $\alpha^{\text{no}}(\phi) = \pi^{\text{no}}(\phi) = 1$ and $p(\phi) = h$ for $\phi > \phi^*$; that is, the unique outcome of the game with no information when $\phi_0 > \phi^*$ is one where all sellers offer h and only the h -buyer accepts such offers.

mimics the ℓ -buyer, his payoff equals $h - \ell$ multiplied by the discounted times the price ℓ is offered. Since $(*)$ is equal to ℓ multiplied by the discounted measure of the times the price ℓ is offered, we have that $(*)$ is equal to $\frac{\ell}{h-\ell} V^x(\phi)$ for all $x \in \{\text{no}, \text{un}, \text{ob}\}$.

The term $(**)$ in equation (19) is equal to h for both $x \in \{\text{no}, \text{un}\}$, since the h -buyer buys with probability one at all times. It is then clear that $W^{\text{un}}(\phi) > W^{\text{no}}(\phi)$ for all $\phi \in (\phi^*, 1)$, because the presence of a signal increase the equilibrium probability of trade for the ℓ -buyer while leaving it unchanged for the h -buyer. When $x = \text{ob}$, on the other hand, for $\mu > \bar{\mu}$ there is a wide range of posteriors where the h -buyer purchases with probability less than one. Additionally, given the slowness of learning due to the rejection mechanism, the ℓ -buyer purchases less often in the public offers case than in the private offers case if μ is large enough. This implies that some transactions that occur in the other cases are not realized in the public offers case. Thus, making prices unobservable may improve social welfare when the signal is very informative.

5 Discussion

There is increasing interest in regulating the availability of information in markets. For example, the European Union, India, Australia, and California have recently passed laws regulating online cookies and other tracking technologies. Similarly, purchases by government entities and listed firms are often subject to transparency regulations.¹⁸ While the main goal of such regulations may not necessarily be to increase market efficiency (they are typically aimed at protecting consumer privacy, preventing corruption, or promoting investor trust), they do affect the incentives of market participants. Our paper provides a canonical setting in which to analyze how changes in the availability of information affect equilibrium welfare and efficiency. We now explain how our results and their implications can be interpreted and used in practice.

Additional information: An important prediction of our model is that buyer welfare and efficiency tend to be greater when more information is available about previous transactions. In the context of online shopping, however, cookies may contain non-transactional information about websites visited by a user, including social media, news websites, and entertainment platforms. This information may reveal individual characteristics such as the user's income, age, gender, or occupation, which sellers may use to determine their price offers. Similarly, in the procurement context, information about the financial po-

¹⁸ The pioneer regulation governing internet privacy was the European Union's General Data Protection Regulation (GDPR), implemented in 2018 and commonly referred to as the Cookie Law. Also, several articles in the Treaty on the Functioning of the European Union (TFEU) concern transparency and openness in decision-making, which are seen as foundational values of the EU.

sition of a public or private entity may be revealed through investments, mandatory public reports, or leaks.

To study the effect of such additional information, one could consider a model with an exogenous stochastic process providing information about the buyer's type, à la Daley and Green (2012).¹⁹ In such a model, our two equilibrium mechanisms would not exist (so, for example, there would not be random discounts), and an increase in signal precision would make the buyer worse off.

In practice, it is reasonable to expect that both types of information are available to sellers. In a combined model, if the exogenous signal were very informative, the buyer's surplus would be close to 0, while if it were very imprecise, the outcome would be close to our predictions. Hence, regulations limiting the overall availability of information would have a double effect. On the one hand, they would limit the sellers' ability to capture information rents, by exogenously limiting their access to the buyer's private information. On the other hand, they would diminish the buyer's incentive to strategically reject offers, so the frequency of discounts would decrease. Thus, regulations may have opposite effects depending on the type of information they address, since they influence both the signaling incentives and the allocation of information rents.

Right to privacy: Some regulations are intended to guarantee the so-called right to privacy, giving users the option to conceal personal information online.²⁰ In our setting, such a possibility could be incorporated by allowing the buyer to decide whether to hide his signal history from each seller. The buyer's decision to hide or reveal his signal history would then convey information about his type. Standard unraveling arguments (à la Milgrom, 1981) would favor the existence of equilibria satisfying our characterization (where the buyer is always believed to be of type h unless he reveals his signal history).²¹

While unraveling suggests that opt-out policies are of limited effectiveness, selective restrictions on certain types of information could be more effective. For example, a regulation could ban cookies from recording non-transactional information (such as demographic data and data on visits to news websites or social media), while allowing users to opt in or out of providing transactional data. Such a policy would protect the user's privacy, limiting the signaling effect of concealing non-transactional data (which

¹⁹ One could further allow for buyers to distort their behavior, e.g. via additional browsing activity, to modify the information provided (similarly to our model). We could model such a possibility as in Dilmé (2019), which studies a dynamic signaling game in which the signal depends on effort instead of type.

²⁰ For example, a major outcome of the EU's Cookie Law is the requirement that websites allow users to opt out of cookies.

²¹ There would also be inefficient equilibria in which disclosure of the history would be perceived as a sign that the buyer is of type h , so that in equilibrium the buyer would not disclose his history. Note that it has been argued that the requirement for users to report their cookie preferences at every website they visit is costly to them, generating "opt-out fatigue", which diminishes the quality of their browsing experience. Johnson et al. (2020) find that in the US, even though users express strong privacy concerns, only a very small fraction opt out of targeted online advertising.

we argue above would tend to benefit sellers rather than buyers), while preserving the user's equilibrium gain from the strategic rejection of high prices.

Buyer sophistication: An important assumption in our model is that buyers are sophisticated. This assumption drives our result that permissive information policies tend to benefit buyers. However, while buyer sophistication may be a reasonable assumption in procurement contexts, it is less clearly so in online markets. In our model, if it were common knowledge that the buyer is not sophisticated, the sellers would charge ℓ for $\phi < \phi^*$ and h for $\phi > \phi^*$. The h -buyer's payoff would then be lower than in our original analysis (in which sellers react to buyer sophistication by lowering prices). In fact, the opposite comparative statics would arise for an unsophisticated buyer: as the signal became more informative, buyer welfare and efficiency would decrease.

These observations suggest that, to increase buyer welfare and efficiency, permissive information policies must be combined with buyer education. It is generally not necessary that all buyers be sophisticated as long as sellers do not directly observe sophistication: as long as the fraction of sophisticated buyers is large enough, sellers will tend to offer low prices to avoid rejection by sophisticated buyers.

Personalized product recommendations: In the context of online shopping, individualized offers often take the form of personalized product recommendations—for example, via advertisements or within-platform promotion of goods. For example, sellers can promote the advertisement of certain goods on platforms such as Google or Amazon as a function of the user's search history.

Personalized product recommendations have effects similar to those of personalized pricing. In their most extreme version, they are equivalent: a seller can offer the same product at different prices (perhaps under different IDs) and strategically promote them depending on the buyer's history. If searching is costly, buyers with high willingness to trade may accept a high price. Similarly, a seller may recommend highly profitable goods to buyers who are likely to have high willingness to trade, while offering better-value goods to other buyers. Therefore, policies aimed at limiting personalized pricing should also address personalized recommendations and advertising, as these can often be used to achieve similar results.

Efficiency: A related point often raised in defense of cookies is that they can improve the user's experience by making internet browsing more efficient. For example, cookies can be used to tailor advertisements to the user's interests, reducing frictions in the market for advertising. (On the other hand, such tailored advertising may also be harmful; for example, it may facilitate political manipulation by third parties.)

To capture these potential benefits, our model could be adapted to allow for some horizontal differentiation. For example, we could allow each seller to offer the buyer one of two types of products, with the assumption that the buyer's relative preference for each product is correlated with his willingness to

trade. The usual tradeoff between efficiency and buyer welfare would then arise: the more information the sellers have, the more efficient trade becomes (as the buyer is offered the goods he prefers) and the more easily the sellers can extract surplus from the buyer. It is then plausible that the h -buyer would sometimes benefit from rejecting his preferred product in order to get the other product at a low price. As in our main model, this would induce sellers to offer low prices. The study of optimal information policies in this setting is left for future research.

As shown by the examples above, the dynamic considerations incorporated into our model help addressing a number of possible effects of policies concerning internet privacy and financial transparency. Such policies affect the information available, thus distorting the incentives for market participants to reveal their private information. Our results imply that allowing more informative signals (e.g., deregulating cookies or promoting transparency) tends to enhance efficiency by limiting the bargaining power of uninformed agents. Nevertheless, some caution is needed: in certain situations, the disclosure of prices or non-transactional information may lower welfare in equilibrium.

References

- BONATTI, A. AND G. CISTERNAS (2020): “Consumer scores and price discrimination,” *The Review of Economic Studies*, 87, 750–791.
- DALEY, B. AND B. GREEN (2012): “Waiting for news in the market for lemons,” *Econometrica*, 80, 1433–1504.
- DILMÉ, F. (2019): “Dynamic quality signaling with hidden actions,” *Games and Economic Behavior*, 113, 116–136.
- FAINGOLD, E. AND Y. SANNIKOV (2011): “Reputation in continuous-time games,” *Econometrica*, 79, 773–876.
- FUDENBERG, D. AND D. K. LEVINE (1989): “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 759–778.
- HART, O. D. AND J. TIROLE (1988): “Contract renegotiation and Coasian dynamics,” *The Review of Economic Studies*, 55, 509–540.
- HÖRNER, J. AND N. VIEILLE (2009): “Public vs. private offers in the market for lemons,” *Econometrica*, 77, 29–69.
- JOHNSON, G. A., S. K. SHRIVER, AND S. DU (2020): “Consumer privacy choice in online advertising: Who opts out and at what cost to industry?” *Marketing Science*, 39, 33–51.
- KAYA, A. AND Q. LIU (2015): “Transparency and price formation,” *Theoretical Economics*, 10, 341–383.
- KAYA, A. AND S. ROY (2022): “Price transparency and market screening,” .
- (2024): “Repeated trading: Transparency and market structure,” *The American Economic Review*, 115, 2388–2435.

- KREPS, D. M. AND R. WILSON (1982): “Reputation and imperfect information,” *Journal of Economic Theory*, 27, 253–279.
- LEE, J. AND Q. LIU (2013): “Gambling reputation: Repeated bargaining with outside options,” *Econometrica*, 81, 1601–1672.
- LIU, Q. (2011): “Information acquisition and reputation dynamics,” *The Review of Economic Studies*, 78, 1400–1425.
- LIU, Q. AND A. SKRZYPACZ (2014): “Limited records and reputation bubbles,” *Journal of Economic Theory*, 151, 2–29.
- MILGROM, P. AND J. ROBERTS (1982): “Predation, reputation, and entry deterrence,” *Journal of Economic Theory*, 27, 280–312.
- MILGROM, P. R. (1981): “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 380–391.
- PEI, H. (2024a): “Reputation effects under short memories,” *Journal of Political Economy*.
- (2024b): “Reputation effects with endogenous records,” .
- STRULOVICI, B. (2017): “Contract negotiation and the Coase conjecture: A strategic foundation for renegotiation-proof contracts,” *Econometrica*, 85, 585–616.
- VILLAS-BOAS, J. M. (2004): “Price cycles in markets with customer recognition,” *RAND Journal of Economics*, 486–501.

A Proofs of the results

A.1 Proofs of results in Section 3

Proof of Lemma 1

Proof. We divide the proof into 3 steps:

Step 1: Preliminary observations. Note that sellers never offer a price below ℓ in equilibrium, since offering ℓ ensures trade. Then, it is without loss of generality to focus on equilibria where the ℓ -seller accepts a price offer if and only if it is equal to ℓ , and where each seller never offers a price lower than ℓ . As a result, $V_\ell(\phi) = 0$ for all $\phi \in (0, 1)$.

Note also that, if the h -buyer accepts a price $\hat{p} > \ell$ for sure, then a seller is willing to offer such a price if and only if $\hat{p} \geq \ell/\phi$. Hence, no price in $(\ell, \ell/\phi)$ is offered in equilibrium.

Finally, it will sometimes be convenient to use the log-likelihood instead of the posterior. For each $\phi \in (0, 1)$ and $z \in \mathbb{R}$, we define

$$\check{z}(\phi) = \log\left(\frac{\phi}{1-\phi}\right) \text{ and } \check{\phi}(z) = \frac{e^z}{1+e^z}. \quad (20)$$

Note that $\check{\phi}(\cdot)$ is the inverse of $\check{z}(\cdot)$. Abusing notation, for some function f of ϕ , we will sometimes use $f(z)$ and $f'(z)$ to denote $f(\check{\phi}(z))$ and $\frac{d}{dz}f(\check{\phi}(z))$, respectively. Note that, for example, $f'(\check{z}(\phi)) = \phi(1-\phi)f'(\phi)$.

Step 2: Monotonicity of V_h . We continue the proof by stating and proving the following result:

Lemma 3. $V_h(\phi) = h - \ell$ for all $\phi \in (0, \phi^*]$ and V_h is strictly decreasing on $(\phi^*, 1)$.

Proof. (i) **Proof that $V_h(\phi) = h - \ell$ for all $\phi \in (0, \phi^*)$.** If only ℓ is offered in $(0, \phi^*)$ then the result follows. Assume then, for the sake of contradiction, that it is optimal for a seller to offer $\hat{p} > \ell$ when the posterior is $\phi \in (0, \phi^*]$. Since \hat{p} has to be at least ℓ/ϕ by Step 1—hence strictly higher than h —, it must be that \hat{p} is accepted for sure (by the argument laid out in the main text after Lemma 1). Then, from equation (9), such price should then satisfy:

$$\ell/\phi \leq \hat{p} = h + \mu/r V_h'(\check{z}(\phi)). \quad (21)$$

Note that, if $|V_h'(\check{z}(\phi))|$ is small enough, the right-hand side of the previous expression is close to h . This implies that, if $|V_h'(\check{z}(\phi))|$ is small enough, sellers offer ℓ for sure (since no price above h is optimal), and $V_h(\phi) = h - \ell$. We conclude that either $V_h'(\phi) = 0$ for all $\phi \in (0, \phi^*)$, in which case $V_h(\phi) = h - \ell$ for all $\phi \in (0, \phi^*]$, or $V_h'(\phi) \neq 0$ for all $\phi \in (0, \phi^*)$.²² In other words, when restricted to $(0, \phi^*)$, either V_h is strictly decreasing, or strictly increasing, or equal to $h - \ell$. It is clear V_h cannot be strictly decreasing, since in this case the h -buyer rejects all prices above h , hence it is strictly optimal for each seller to offer ℓ for all $\phi \in (0, \phi^*)$ and $V_h(\phi) = h - \ell$. Assume, for the sake of contradiction, that V_h is strictly increasing. Equation (21) implies that, in this case, $\lim_{z \rightarrow -\infty} V_h'(z) = +\infty$ and hence $\lim_{z \rightarrow -\infty} V_h(z) = -\infty$, which is a contradiction. Then, the only possibility is that $V_h(\phi)$ is equal to $h - \ell$ for all $\phi \in (0, \phi^*]$.

(ii) **Proof that V_h is strictly decreasing on $(\phi^*, 1)$.** Assume first $V_h'(\phi_1) = 0$ for some $\phi_1 \in (\phi^*, 1)$. This implies that $\alpha_h(\phi_1, \hat{p}) = 1$ for all $\hat{p} < h$, and $\alpha_h(\phi_1, \hat{p}) = 0$ for all $\hat{p} > h$. Hence, at posterior ϕ_1 , the seller offers h for sure in equilibrium, since $\phi h > \ell$ when $\phi > \phi^*$. Equation (4) then becomes

$$r V_h(\phi_1) = \frac{1}{2} \mu (1 - \phi_1)^2 \phi_1^2 V_h''(\phi_1).$$

We then have that ϕ_1 is a minimizer of V_h . This implies that V_h is strictly increasing on $(\phi_1, 1)$, so all prices offered on $(\phi_1, 1)$ are higher than h (by the argument used to obtain equation (9)). But

²² The observation implies that V_h' cannot continuously “approach” 0, since a small value of $V_h'(\phi)$ implies $V_h(\phi) = h - \ell$ when $\phi < \phi^*$.

then this implies that $\lim_{\phi \nearrow 1} V_h(\phi) \leq 0$, and therefore $V_h(\phi) < 0$ for some ϕ , which is a contradiction. We conclude that $V_h'(\phi) \neq 0$ for all $\phi \in (\phi^*, 1)$.

Since V_h cannot be strictly increasing on $(\phi^*, 1)$ (because $V_h(\phi^*) = h - \ell$ and equilibrium offers are never lower than ℓ) and $V_h'(\phi) \neq 0$ for all $\phi \in (\phi^*, 1)$, we have that the V_h must be strictly decreasing on $(\phi^*, 1)$. \square

(End of the proof of Lemma 3. Proof of Lemma 1 continues.)

Step 3: Proof of the result. The argument in the main text implies that either the seller offers ℓ , or $p(\phi)$ satisfying equation (9), or randomizes between them. An immediate implication of Lemma 3 is that ℓ is offered with probability one when $\phi \in (0, \phi^*]$. Also, since V_h is strictly decreasing on $(\phi^*, 1)$, we have that ℓ is never offered with probability one at posteriors in $(\phi^*, 1)$ (note that if ℓ was offered with probability one at some posterior $\phi \in (\phi^*, 1)$, then the posterior would remain equal to ϕ at all times for $\phi_0 = \phi$, so we would have $V_h(\phi) = h - \ell = V_h(\phi^*)$). \square

Proof of Propositions 1 and 2

Proof. We prove Propositions 1 and 2 together. We divide the proof into ten steps:

Step 1: Preliminary derivations. We begin the proof by providing a useful equation. The arguments in the main text (following Lemma 1) show that, for each posterior $\phi \in (\phi^*, 1)$, the buyer is indifferent between accepting $p(\phi)$ or not (that is, equation (7) holds for $\hat{p} = p(\phi)$). This implies that the buyer's continuation value can be computed as if he did reject $p(\phi)$ for all posteriors ϕ . As a result, the continuation value of the buyer satisfies the following equation:

$$\begin{aligned} rV(\phi) &= (1 - \pi(\phi))(h - \ell) - \mu(1 - \phi)\phi^2\pi(\phi)\alpha(\phi)^2V'(\phi) \\ &\quad + \frac{1}{2}\mu(1 - \phi)^2\phi^2\pi(\phi)\alpha(\phi)^2V''(\phi). \end{aligned} \tag{22}$$

The previous expression is convenient as it does not depend on $p(\phi)$.

Step 2: Preliminary results on continuity. We continue by providing a result on the continuity of $\alpha(\cdot)$ and $p(\cdot)$, and the limits of $V(\phi)$ as ϕ tends to ϕ^* and 1.

Lemma 4. *Both $\alpha(\cdot)$ and $p(\cdot)$ are continuous on $(\phi^*, 1)$. Furthermore, $\lim_{\phi \searrow \phi^*} V(\phi) = h - \ell$ and $\lim_{\phi \nearrow 1} V(\phi) = 0$.*

Proof. The continuity of $\alpha(\cdot)$ and $p(\cdot)$ on $(\phi^*, 1)$ is immediately implied by equations (8) and (9). That $\lim_{\phi \searrow \phi^*} V(\phi) = h - \ell$ follows from continuity of the continuation value and Lemma 3.

We finally prove that $\lim_{\phi \nearrow 1} V(\phi) = 0$. Recall that, by Lemma 3, V is strictly decreasing on $(\phi^*, 1)$. We assume, for the sake of contradiction, that $\lim_{\phi \nearrow 1} V(\phi) > 0$. Since it is optimal for the h -buyer to follow the ℓ -buyer's strategy (that is, only accepting offers equal to ℓ), there must exist some increasing sequence $(\phi_n)_n$ converging to 1 such that $(\pi(\phi_n))_n$ is convergent and $\lim_{n \rightarrow \infty} \pi(\phi_n) < 1$.²³ We can write equation (9) using log-likelihoods (see equation (20)) as

$$p(\phi) = \max\{h/2, h + \mu/r V'(\check{z}(\phi))\}.$$

Hence, since $\lim_{\phi \nearrow 1} V'(\check{z}(\phi)) = 0$ (because V is strictly decreasing and bounded below by 0), we have that $\lim_{\phi \nearrow 1} p(\phi) = h$. This implies that, if ϕ is close enough to 1, $p(\phi) > \max\{\ell/\phi, h/2\}$, hence $\pi(\phi) = 1$ (by (12)). As a result, $\lim_{\phi \nearrow 1} V(\phi) = 0$. \square

(End of the proof of Lemma 4. Proof of Propositions 1 and 2 continues.)

Step 3: Preliminary results on regimes. We now present a result providing the equations for each of four possible types of regimes on $(\phi^*, 1)$.

Lemma 5. *The following statements follow for all $\phi_1, \phi_2 \in [\phi^*, 1]$ with $\phi_1 < \phi_2$:*

(i) *If $p(\phi) = \phi/\ell$ for all $\phi \in (\phi_1, \phi_2)$ then $\phi_2 \leq 2\phi^*$, $\alpha(\phi) = 1$, $\pi(\phi) \in (0, 1)$,*

$$\pi(\phi) = \frac{2\phi(h-\ell-V(\phi))}{\phi(h-\ell)+2\ell(1-\phi)}, \quad (23)$$

and

$$\pi'(\phi) = \frac{2r(\phi-\phi^*)+2(1-\phi)\phi^*\mu\pi(\phi)}{(1-\phi)\phi\mu(\phi(1-\phi^*)+2(1-\phi)\phi^*)} > 0. \quad (24)$$

(ii) *If $p(\phi) = h/2$ for all $\phi \in (\phi_1, \phi_2)$ then $\phi_1 \geq 2\phi^*$ (hence $\phi^* < 1/2$), $\alpha(\phi) \in [2\phi^*/\phi, 1)$, and $V(\phi) > h/4$. Also,*

(a) *If $\alpha(\phi) = 2\phi^*/\phi$ for all $\phi \in (\phi_1, \phi_2)$, then*

$$\pi(\phi) = \frac{2\phi(h-\ell-V(\phi))}{2h-3\ell} \quad (25)$$

²³ If no such sequence exists, then the fraction of time where ℓ is offered during the time the posterior is in $(\bar{\phi}, 1)$ shrinks to 0 as $\bar{\phi} \rightarrow 1$. This implies that the payoff the h -seller obtains by mimicking the ℓ -seller converges to 0 as ϕ gets close 1.

and

$$\pi'(\phi) = \frac{r}{2\mu(1-\phi)(2-3\phi^*)\phi^*} > 0. \quad (26)$$

(b) If $\alpha(\phi) \in (2\phi^*/\phi, 1)$ for all $\phi \in (\phi_1, \phi_2)$, then $\pi(\phi) = 1$,

$$\alpha(\hat{p}) = -\frac{rh}{2\mu(1-\phi)\phi V'(\phi)}, \quad (27)$$

and

$$\alpha'(\phi) = \frac{4V(\phi) - h\alpha(\phi)}{(1-\phi)\phi h} > 0. \quad (28)$$

(iii) If $p(\phi) > \max\{h/2, \ell/\phi\}$ for all $\phi \in (\phi_1, \phi_2)$ then $\alpha(\phi) = 1$, $\pi(\phi) = 1$, and

$$p'(\phi) = \frac{2V(\phi) - h + p(\phi)}{(1-\phi)\phi} > 0. \quad (29)$$

Proof. (i) The fact that $\phi_2 \leq 2\phi^*$ follows from the fact that $p(\phi) \geq h/2$ for all ϕ (by equation (9)). The fact that $\alpha(\phi) = 1$ follows from equation (10). Equations (23) and (24) follow from equations (9) (with $p(\phi) = \ell/\phi$) and (22) (with $\alpha(\phi) = 1$). Finally, $\pi(\phi) \in (0, 1)$ because $\pi(\phi) \in [0, 1]$ and π is strictly increasing by equation (24).

(ii) That $\phi_1 \geq 2\phi^*$ follows from equation (11). That $\alpha(\phi) \geq 2\phi^*/\phi$ follows from the definition of ϕ^* and the optimality condition for the sellers requiring that $\phi \alpha(\phi) p(\phi) \geq \ell$.

That $\alpha(\phi) < 1$ follows from the fact that, if $\alpha(\phi) > 2\phi^*/\phi$, then the first equality in equation (28) holds (from equations (22) and (27) with $\pi(\phi) = 1$), so $\alpha(\phi) = 1$ and $p(\phi) = h/2$ only if $V(\phi) = h/4$. That $\alpha(\phi) < 1$ for all $\phi \in (\phi_1, \phi_2)$ follows from the fact that V is strictly increasing on $(\phi^*, 1)$ (by Lemma 3). Also:

(a) Equation (26) follows from equations (22) and equation (7) (with $\alpha(\phi) = 2\phi^*/\phi$).

(b) That $\pi(\phi) = 1$ follows because the seller strictly prefers offering $h/2$ than offering ℓ (since $\ell < \phi \alpha(\phi) p(\phi)$). Equation (27) follows from equation (8), and equation (28) follows from equations (22) and (27) (with $\pi(\phi) = 1$).

To prove that $V(\phi) > h/4$, recall the end of the proof of Lemma 4. It shows that there exists some $\bar{\phi} < 1$ such that $\alpha(\phi) = 1$ for all $\phi \in [\bar{\phi}, 1)$. This implies that, if $\alpha(\phi) = 1$ for all $\phi \in (\phi_1, \phi_2)$, then there must be some $\bar{\phi} < 1$ such that $\alpha'(\bar{\phi}) \geq 0$ and so $V(\bar{\phi}) \geq h/4$. Since V is strictly decreasing on $(\phi^*, 1)$ by Lemma 3, we have $V(\phi) > h/4$ for all $\phi \in (\phi_1, \phi_2)$.

(iii) That $\alpha(\phi)=1$ follows from equation (10). That $\pi(\phi)=1$ follows from equation (12). Equation (29) follows from differentiating equation (9) (since the max operator on its right-hand side is larger than $h/2$) and from using equation (22). That $p'(\phi)>0$ follows from the fact that, from equations (9) and (29), we have that $p''(\phi)<0$ when $p'(\phi)=0$, but we know that $p(\phi)<h$ and $\lim_{\phi' \nearrow 1} p(\phi')=1$. \square

(End of the proof of Lemma 5. Proof of Propositions 1 and 2 continues.)

Step 4: Proof that $p(\phi)=\ell/\phi$ and $\alpha(\phi)=1$ for ϕ close to ϕ^* . Take a sequence $(\phi_n)_n$ strictly decreasing toward ϕ^* such that $\phi_n \in (\phi^*, \min\{1, 2\phi^*\})$ for all n . From equation (11), we have that $p(\phi_n) \geq \ell/\phi_n > h/2$ for all n , and so $p(\phi_n) \rightarrow h$. By Lemma 5, we have that $\alpha(\phi_n)=1$ for all n . Recall that, from Lemma 4, $V(\phi_n)$ converges to $h-\ell$. Assume, for the sake of contradiction and taking a subsequence if necessary, that $p(\phi_n) > \ell/\phi_n > h/2$ for all n . Lemma 5 implies that $p(\phi)$ is increasing when it is strictly larger than $\max\{\ell/\phi_n, h/2\}$, hence this implies that $p(\phi) > \max\{\ell/\phi_n, h/2\}$ for all $\phi \in (\phi^*, 1)$. From equation (29) it is clear that there must be some posterior $\phi' > \phi^*$ such that $p(\phi') > h$, which contradicts Lemma 1.

Step 5: Definition of ϕ^\dagger . From the previous step, there must be some maximal ϕ^\dagger such that $p(\phi)=\ell/\phi$ for all $\phi \in (\phi^*, \phi^\dagger)$. Take some $\phi \in (\phi^*, \phi^\dagger)$; that is, we have that $p(\phi)=\ell/\phi$ and $\alpha(\phi)=1$. From equation (9) we have

$$\ell/\phi = p(\phi) = h + r^{-1} \mu (1-\phi) \phi V'(\phi). \quad (30)$$

Additionally, by Lemma 5, we have $\pi(\phi) \in (0, 1)$. We can solve (30) for V (with boundary condition $V(\phi^*)=h-\ell$), and obtain

$$V(\phi) = h - \ell + \frac{r(\phi h - \ell)}{\mu \phi} + \frac{r(h - \ell)}{\mu} \log\left(\frac{1-\phi}{\phi} / \frac{1-\phi^*}{\phi^*}\right). \quad (31)$$

Note that we have $V'(\phi^*)=0$, so V' is continuous at ϕ^* . Then, from equation (22), we have

$$\pi(\phi) = \frac{2r}{(\phi(h-3\ell)+2\ell)\mu} \left(\ell - \phi h - \phi (h - \ell) \log\left(\frac{1-\phi}{\phi} / \frac{1-\phi^*}{\phi^*}\right) \right). \quad (32)$$

Since the right-hand side of equation (31) tends to $-\infty$ as $\phi \rightarrow 1$, it must be that $\phi^\dagger < 1$.

Step 6: Preliminaries for less informative signal. We first focus on the case $\phi^\dagger < 2\phi^*$. We want to show that, for all $\phi > \phi^\dagger$, we have $p(\phi) > \max\{h/2, \ell/\phi\}$, and hence $\alpha(\phi)=1$, $\pi(\phi)=1$, and equation (29) holds (by Lemma 5).

We first argue that there is no $\phi \in (\phi^\dagger, \min\{2\phi^*, 1\})$ where $p(\phi)=\ell/\phi$. Assume, for the sake of

contradiction, that there is an interval (ϕ_1, ϕ_2) with $\phi^\dagger \leq \phi_1 < \phi_2 \leq 2\phi^*$ such that $p(\phi) > \ell/\phi$ for all $\phi \in (\phi_1, \phi_2)$ and $p(\phi_2) = \ell/\phi_2$. Since $p(\phi)$ approaches ℓ/ϕ_2 as $\phi \rightarrow \phi_2$ from above the curve ℓ/ϕ , it must be that $\lim_{\phi \nearrow \phi_2} p'(\phi) \leq -\ell/\phi_2^2$ which, by equation (29), implies

$$V(\phi_2) \leq \frac{1}{2}(h + \ell - 2\ell/\phi_2) < \frac{1}{2}(\ell - h) < 0,$$

which is a contradiction. Hence, there is no $\phi \in (\phi^\dagger, \min\{2\phi^*, 1\})$ where $p(\phi) = \ell/\phi$.

We now argue that there is no $\phi \in (\min\{1, 2\phi^*\}, 1)$ such that $p(\phi) = h/2$. If $\phi^* \geq 1/2$ the result is clear, so assume that $\phi^* < 1/2$. Assume, for the sake of contradiction, that there is an interval (ϕ_1, ϕ_2) with $2\phi^* \leq \phi_1 < \phi_2 < 1$ such that $p(\phi) > h/2$ for all $\phi \in (\phi_1, \phi_2)$ and $p(\phi_2) = h/2$. It is clear that $p(\phi) > \max\{\ell/\phi, h/2\}$ for all $\phi \in (\phi_1, \phi_2)$ (since $h/2 > \ell/\phi$ for $\phi > 2\phi^*$). This implies, by Lemma 5, that p is increasing on (ϕ_1, ϕ_2) , hence it is not possible that $p(\phi_2) = h/2$ (p is continuous by Lemma 4).

We then have proven that $p(\phi) > \max\{h/2, \ell/\phi\}$ and $\alpha(\phi) = \pi(\phi) = 1$ for all $\phi > \phi^\dagger$; hence, by Lemma 5, equation (29) holds. This implies that equation (22) can be written, for all $\phi \in (\phi^\dagger, 1)$, as

$$\begin{aligned} rV(\phi) &= (1 - \pi(\phi))(h - \ell) - \mu(1 - \phi)\phi^2 \pi(\phi)^2 \alpha(\phi)^2 V'(\phi) \\ &\quad + \frac{1}{2}\mu(1 - \phi)^2 \phi^2 \pi(\phi)^2 \alpha(\phi)^2 V''(\phi). \end{aligned} \quad (33)$$

The solution to this equation is

$$V(\phi) = C_1 \left(\frac{1 - \phi}{\phi}\right)^\kappa + C_2 \left(\frac{1 - \phi}{\phi}\right)^{-1 - \kappa}, \quad (34)$$

where

$$\kappa := \frac{1}{2}(\sqrt{1 + 8r/\mu} - 1) > 0, \quad (35)$$

and C_1 and C_2 are integration constants. Using that, by Lemma 4, we have that $\lim_{\phi \nearrow 1} V(\phi) = 0$, and so $C_2 = 0$. We can use the continuity of V' at ϕ^\dagger and equations (31) and (34) with $C_2 = 0$ to obtain the value of C_1 as a function of ϕ^\dagger , so we have that

$$V(\phi) = \frac{r}{\mu} \frac{\phi^\dagger - \phi^*}{\phi^\dagger} \left(\frac{\phi}{1 - \phi} / \frac{\phi^\dagger}{1 - \phi^\dagger}\right)^{-\kappa} h \quad (36)$$

for all $\phi \in (\phi^\dagger, 1)$. Finally, the value of ϕ^\dagger is obtained using that V is continuous at ϕ^\dagger . This requirement can be written as

$$-\frac{(\kappa - 1)r(\phi^\dagger - \phi^*)}{\phi^{\dagger\kappa}} + r(1 - \phi^*) \log\left(\frac{\phi^\dagger}{1 - \phi^\dagger} / \frac{\phi^*}{1 - \phi^*}\right) - \mu(1 - \phi^*) = 0. \quad (37)$$

The left-hand side of the previous expression is $-\mu(1-\phi^*)$ when $\phi^\dagger = \phi^*$, and tends to $+\infty$ when $\phi^\dagger \rightarrow 1$, so a solution exists. Furthermore, the derivative of the left-hand side of the previous equation with respect to ϕ^\dagger is

$$\frac{r(\kappa(\phi^\dagger - \phi^*) + (1 - \phi^\dagger)\phi^*)}{\kappa(1 - \phi^\dagger)(\phi^\dagger)^2} > 0.$$

Hence, there is exactly one value of ϕ^\dagger solving equation (37). The derivative of the left-hand side of expression (36) with respect to μ is

$$-\frac{2r^2(\phi^\dagger - \phi^*)}{\mu^2 \phi^\dagger (1 + 2\kappa)\kappa^2} - (1 - \phi^*) < 0.$$

It is then clear that ϕ^\dagger is increasing in μ .

Step 7: Definition of $\bar{\mu}$. We now claim that there is some value $\bar{\mu} \in (0, +\infty]$ such that a solution ϕ^\dagger strictly smaller than $2\phi^*$ for equation (37) exists if and only if $\mu < \bar{\mu}$. The result is obviously true if $\phi^* \geq 1/2$, since then $\bar{\mu} = +\infty$ (by the arguments in Step 6). Assume then that $\phi^* < 1/2$.

Note that, differentiating the left-hand side of expression (37) two times with respect to μ (recall that κ depends on μ , see equation (35)), we obtain

$$-\frac{7}{8} + \phi^* + \frac{\mu + 4r}{8\mu(1 + 2\kappa)}.$$

Using the value of κ (from (35)) it is easily seen that the previous expression is negative for all $\phi^* < 1/2$. Furthermore, the left-hand side of expression (37) is negative when $\phi^\dagger = \phi^*$, and tends to $+\infty$ when $\phi^\dagger \rightarrow 1$. It is then clear that there is only one value of μ for which $\phi^\dagger = 2\phi^*$. That is, using Step 6, we have proven that, for all $\mu \leq \bar{\mu}$, the unique equilibrium is as described in the state of Proposition 1, while for all $\mu > \bar{\mu}$ there is no equilibrium of this form.²⁴

Step 8: Preliminaries for the more informative signal. Assume for the rest of the proof that $\mu > \bar{\mu}$. In this case, $\phi^\dagger = 2\phi^*$ (where ϕ^\dagger is defined in Step 5 as the maximal such that $p(\phi) = \ell/\phi$ for all $\phi \in (\phi^*, \phi^\dagger)$). Since $\alpha(\cdot)$ is continuous on $(\phi^*, 1)$ (by Lemma 4), an implication of Lemma 5 is that all equilibria satisfy the characterization provided in Proposition 2.²⁵

Then, to show existence and uniqueness of equilibria, we have prove the existence and uniqueness

²⁴ Indeed, as $\mu \rightarrow 0$ we have $\phi^\dagger \rightarrow \phi^*$, so the value of ϕ^\dagger is increasing in μ at $\phi^\dagger = 2\phi^*$. (Proposition 3 shows that the unique ϕ^\dagger solving equation (37) is increasing in μ .)

²⁵ Indeed, recall that α is continuous by Lemma 4. By Lemma 5, it is decreasing only if it is equal to $2\phi^*/\phi$. Hence, if an equilibrium exists for $\mu > \bar{\mu}$, there must be some $\hat{\phi}^\dagger$ and $\hat{\phi}^{\dagger\dagger}$ such that $\alpha(\phi) = 2\phi^*/\phi$ for all $\phi \in (2\phi^*, \hat{\phi}^\dagger)$ (case 2(a) in Lemma 5), then satisfies equation (27) in $(\hat{\phi}^\dagger, \hat{\phi}^{\dagger\dagger})$ (where it is strictly increasing), and it is equal to 1 in $(\hat{\phi}^{\dagger\dagger}, 1)$.

thresholds $\hat{\phi}^\dagger$ and $\hat{\phi}^{\dagger\dagger}$ such that the implied continuation value satisfies the smooth pasting conditions and such that $\alpha(\hat{\phi}^{\dagger\dagger})=1$.

To prove existence of an equilibrium, we construct the continuation value of the h -buyer. To do so, we use Lemma 5 to determine the equations governing V for the different regions of beliefs, for some given values of $\hat{\phi}^\dagger$ and $\hat{\phi}^{\dagger\dagger}$.

(i) Region $(0, \phi^*)$: In this region we have $V(\phi)=h-\ell$ (by Lemma 1).

(ii) Region $(\phi^*, 2\phi^*)$: In this region, equation (31) holds.

(iii) Region $(2\phi^*, \hat{\phi}^\dagger)$: Imposing the smooth pasting conditions at $2\phi^*$, we obtain that, in this region,

$$V(\phi)=h-\ell+\left(\frac{h}{2}+\frac{\ell}{4}\log\left(\frac{1-\phi}{1-2\phi^*}\right)-(h-\ell)\log\left(\frac{2-2\phi^*}{1-2\phi^*}\right)\right)\frac{r}{\mu}. \quad (38)$$

(iv) Region $(\hat{\phi}^\dagger, \hat{\phi}^{\dagger\dagger})$: In this region, V follows equation (22) with $\pi(\phi)=1$. Using (27), we obtain

$$V(\phi)=\frac{1}{2\sqrt{\mu/r}}hH\left((1-\phi)^{-1}c_1+c_2\right)$$

for some $c_1 < 0$ (so V is decreasing) and $c_2 \in \mathbb{R}$; where H is the inverse of the integral of the Gaussian distribution multiplied by 2, that is,

$$H^{-1}(x):=\frac{2}{\sqrt{\pi}}\int_0^x e^{-x'^2} dx'.$$

Note that the domain and codomain of H are, respectively, $(-1, 1)$ and \mathbb{R} , and also that H is strictly increasing and strictly convex on $(0, 1)$.

(v) Region $(\hat{\phi}^{\dagger\dagger}, 1)$: In this region, V follows equation (34) for some $C_1 > 0$ and $C_2 = 0$.

We have 5 unknown variables to be determined: $\hat{\phi}^\dagger$, $\hat{\phi}^{\dagger\dagger}$, c_1 , c_2 and C_1 . To do so, we have two smooth pasting conditions at $\hat{\phi}^\dagger$, two other smooth pasting conditions at $\hat{\phi}^{\dagger\dagger}$, and the requirement that $\alpha(\hat{\phi}^{\dagger\dagger})=1$. We have then as many unknown variables as conditions.

Using the smooth pasting conditions at $\hat{\phi}^{\dagger\dagger}$ and that $\alpha(\hat{\phi}^{\dagger\dagger})=1$, we obtain

$$C_1=\frac{rh}{2\kappa\mu}\left(\frac{\hat{\phi}^{\dagger\dagger}}{1-\hat{\phi}^{\dagger\dagger}}\right)^\kappa$$

(where recall that κ is defined in (35)) and also

$$H\left((1-\hat{\phi}^{\dagger\dagger})^{-1}c_1+c_2\right)=\left(\kappa\sqrt{\mu/r}\right)^{-1} \quad (39)$$

and

$$H'((1-\hat{\phi}^{\dagger\dagger})^{-1}c_1+c_2)=-\frac{1-\hat{\phi}^{\dagger\dagger}}{\hat{\phi}^{\dagger\dagger}c_1\sqrt{\mu/r}}. \quad (40)$$

We define the variables $c^\dagger := (1-\hat{\phi}^\dagger)^{-1}c_1+c_2$ and $c^{\dagger\dagger} := (1-\hat{\phi}^{\dagger\dagger})^{-1}c_1+c_2$. Since V is positive and decreasing and H is increasing, it must be that $c^\dagger > c^{\dagger\dagger} > 0$. From the two equations (39) and (40), and from the boundary conditions at $\hat{\phi}^\dagger$ we obtain the following four equations:

(i) The first equation determines the value of $c^{\dagger\dagger}$:

$$c^{\dagger\dagger} = H^{-1}\left(\frac{1}{\kappa\sqrt{\mu/r}}\right) = \frac{2}{\sqrt{\pi}} \int_0^{\frac{r^{1/2}}{\kappa\mu^{1/2}}} e^{-x^2} dx. \quad (41)$$

(ii) The second and third equations express $\hat{\phi}^\dagger$ and $\hat{\phi}^{\dagger\dagger}$ as functions of c^\dagger and $c^{\dagger\dagger}$:

$$\begin{aligned} \hat{\phi}^\dagger &= 2\phi^* \frac{H'(c^\dagger)}{H'(c^{\dagger\dagger})} (1 - \sqrt{\mu/r}(c^\dagger - c^{\dagger\dagger})H'(c^{\dagger\dagger})), \\ \hat{\phi}^{\dagger\dagger} &= \frac{2\phi^* H'(c^\dagger)}{H'(c^{\dagger\dagger})(1 + 2\phi^* \sqrt{\mu/r}(c^\dagger - c^{\dagger\dagger})H'(c^\dagger))}. \end{aligned}$$

It is not difficult to see that, as long as $c^\dagger > c^{\dagger\dagger}$, we have $\hat{\phi}^\dagger < \hat{\phi}^{\dagger\dagger} < 1$.

(iii) The only value left to determine is c^\dagger . This is obtained by solving the fourth equation

$$\begin{aligned} 0 &= r \log\left((1-2\phi^*)\left(1-2\phi^* H'(c^\dagger) \overbrace{\left(\frac{1}{H'(c^{\dagger\dagger})} - \frac{\sqrt{\mu}(c^\dagger - c^{\dagger\dagger})}{\sqrt{r}}\right)}^{(*)}\right)\right) \\ &\quad + 2\phi^* \left(2(1-\phi^*)\mu - 2(1-\phi^*)r \log\left(1 + \frac{1}{1-2\phi^*}\right) - \sqrt{\mu}\sqrt{r}H(c^\dagger) + r\right). \end{aligned} \quad (42)$$

To prove existence of an equilibrium, we have to prove that equation (42) has a solution for c^\dagger in $(c^{\dagger\dagger}, 1)$, and to prove uniqueness of an equilibrium that such solution is unique.

Step 9: Existence of an equilibrium for the more informative signal. We first note that if c^\dagger is replaced by $c^{\dagger\dagger}$ in equation (42), the resulting equation is equivalent to equation (37) with $\phi^\dagger = 2\phi^*$. This result is intuitive: when $\mu = \bar{\mu}$ (hence equation (37) with $\phi^\dagger = 2\phi^*$ holds), we have $c^\dagger = c^{\dagger\dagger}$, which implies that $\hat{\phi}^\dagger = \hat{\phi}^{\dagger\dagger} = 2\phi^*$. Since in this part of the proof we assume that $\mu > \bar{\mu}$ implies that the right-hand side of equation (42) is positive when c^\dagger is replaced by $c^{\dagger\dagger}$.

It is easy to see that the term $(*)$ in equation (42) is positive when c^\dagger is replaced by $c^{\dagger\dagger}$. Hence, since $H'(c^\dagger)$ tends to $+\infty$ when c^\dagger tends to 1, we have that there exists a value \bar{c}^\dagger such that the right-hand side of equation (42) tends to $-\infty$ as $c^\dagger \nearrow \bar{c}^\dagger$.²⁶ Then, continuity proves the existence of an equilibrium.

²⁶ Indeed, the term inside the logarithm on the first line of equation (42) tends to 0 as $c^\dagger \nearrow \bar{c}^\dagger$.

Step 10: Uniqueness of an equilibrium for the more informative signal. It is only left to prove that the right-hand side of equation (42) is strictly decreasing on $(c^{\dagger\dagger}, \bar{c}^{\dagger})$. To do so, define $\tilde{c}^{\dagger} := H(c^{\dagger})$. Then, the derivative of the right-hand side of equation (42) with respect to \tilde{c}^{\dagger} is

$$0 = -2\phi^* r H''(c^{\dagger}) \overbrace{\frac{1}{H'(c^{\dagger\dagger})} - \frac{\sqrt{\mu}(c^{\dagger} - c^{\dagger\dagger})}{\sqrt{r}} + \frac{\sqrt{\mu} H'(c^{\dagger})}{\sqrt{r} H''(c^{\dagger})}}^{(**)} - 2\phi^* \sqrt{\mu} \sqrt{r}. \quad (43)$$

The derivative of $(**)$ with respect to c^{\dagger} is $\sqrt{\mu}/(2\sqrt{r}H(c^{\dagger})) > 0$, and its value at $c^{\dagger} = c^{\dagger\dagger}$ is

$$\frac{e^{-\frac{r}{\kappa^2 \mu}} (4r + \mu - \sqrt{\mu} \sqrt{8r + \mu})}{2\sqrt{\pi} r},$$

which is positive. Hence, the term $(**)$ is positive. As a result, the derivative of the right-hand side of equation (42) with respect to c^{\dagger} is negative. We conclude that there is a unique pair of thresholds $(\hat{\phi}^{\dagger}, \hat{\phi}^{\dagger\dagger})$, with $2\phi^* < \hat{\phi}^{\dagger} < \hat{\phi}^{\dagger\dagger} < 1$ such that V satisfies all boundary conditions. Then, when $\mu > \bar{\mu}$, a unique equilibrium (which satisfies the characterization of Proposition 2) exists. \square

A.2 Proofs of results in Section 4

Proof of Proposition 3

Proof. The proof is divided into four steps.

Step 1. We first show that V is increasing in μ on $(0, \bar{\mu})$. We divide this step of the proof into two sub-steps:

- (i) We first prove that ϕ^{\dagger} is increasing in μ . To do so, recall from the proof of Propositions 1 and 2 that ϕ^{\dagger} is the unique solution to (37), which we denote ϕ_{μ}^{\dagger} in this proof.

We note that the derivative of the left-hand side of expression (37) with respect to μ (recall that κ depends on μ , see equation (35)) is equal to

$$\frac{\mu(1+\kappa)+2r}{2\mu(1+2\kappa)} \frac{\phi^{\dagger} - \phi^*}{\phi^*} - (1 - \phi^*).$$

Such expression is strictly increasing in ϕ^{\dagger} and negative for all $\phi^{\dagger} \in (\phi^*, \bar{\phi}^{\dagger})$, where

$$\bar{\phi}^{\dagger} := \min \left\{ 1, \phi^* \left(1 - \frac{2(1+2\kappa)\mu(1-\phi^*)}{2r+(1+\kappa)\mu} \right)^{-1} \right\}.$$

If $\bar{\phi}^{\dagger} = 1$ then we have that the left-hand side of expression (37) is decreasing in μ , and hence ϕ^{\dagger} is increasing in μ . Assume then that $\bar{\phi}^{\dagger} < 1$. When we evaluate the left-hand side of expression (37)

for $\phi^\dagger = \overline{\phi^\dagger}$ we obtain

$$-(1 - \phi^\dagger) \left(\mu + \frac{2(\kappa-1)(1+2\kappa)r\mu}{2\kappa r + \kappa(1+\kappa)\mu} + r \log \left(\frac{(1+3\kappa)\mu - 2r}{(1+\kappa)\mu + 2r} \right) \right).$$

Using the definition of κ (see equation (35)), it is easy to see that the previous expression is positive. This implies that $\overline{\phi^\dagger} > \phi^\dagger$. As a result, we have that the derivative of the left-hand side of equation (37) at ϕ_μ^\dagger is negative. Recalling that the left-hand side of equation (37) is strictly increasing in ϕ^\dagger (see the proof of Propositions 1 and 2), we have that ϕ_μ^\dagger is increasing in μ on $(0, \overline{\mu})$.

(ii) Let $z^* := \check{z}(\phi^*)$ and $z^\dagger := \check{z}(\phi^\dagger)$ (recall the definition of \check{z} in (20)). Note that ϕ^\dagger for all $z \in [z^\dagger, \infty)$ we have²⁷

$$\frac{V(z^\dagger)}{-V'(z^\dagger)} \begin{cases} > \kappa^{-1} & \text{if } z \in (z^*, z^\dagger). \\ = \kappa^{-1} & \text{if } z \in [z^\dagger, +\infty). \end{cases} \quad (44)$$

Take now two values $\mu_1 < \mu_2 < \overline{\mu}$. By part 1, we have that $z_{\mu_1}^\dagger < z_{\mu_2}^\dagger$, where the subindexes μ_1 and μ_2 indicate equilibrium values corresponding to each value of μ . From equations (9) (with $p(\phi) = \ell/\phi$) and (31), we have $V_{\mu_1}(z) < V_{\mu_2}(z)$ and $-V'_{\mu_1}(z) > -V'_{\mu_2}(z)$ for all $z \in (z^*, z_{\mu_1}^\dagger]$. Assume, for a contradiction, that there is some $z > z_{\mu_1}^\dagger$ such that $V_{\mu_1}(z) = V_{\mu_2}(z)$. There must then be some $\hat{z} > z_{\mu_1}^\dagger$ such that $V_{\mu_1}(\hat{z}) = V_{\mu_2}(\hat{z})$ and $-V'_{\mu_1}(\hat{z}) \leq -V'_{\mu_2}(\hat{z})$. But then,

$$\frac{V_{\mu_2}(\hat{z})}{-V'_{\mu_2}(\hat{z})} \leq \frac{V_{\mu_1}(\hat{z})}{-V'_{\mu_1}(\hat{z})} = \kappa_{\mu_1}^{-1} < \kappa_{\mu_2}^{-1} \leq \frac{V_{\mu_2}(\hat{z})}{-V'_{\mu_2}(\hat{z})}.$$

This is a contradiction. We then showed that $V_{\mu_1}(\phi) < V_{\mu_2}(\phi)$ for all $\phi \in (\phi^*, 1)$.

Step 2. That $\frac{d}{d\mu} V(\phi) > 0$ when $\mu > \overline{\mu}$ and $\phi \in (\phi^*, \hat{\phi}^\dagger)$ follows immediately from differentiating the right-hand side of expressions (31) (which holds when $\phi \in (\phi^*, 2\phi^*)$) and (38) (which holds when $\phi \in (2\phi^*, \hat{\phi}^\dagger)$).

Step 3. We now show that $\lim_{\mu \rightarrow \infty} V(\phi) = h - \ell$ for all $\phi \in (\phi^*, 1)$. There are two cases:

- (i) Assume first $\overline{\mu} = +\infty$ (that is, if $\phi^* \geq 1/2$). Take a sequence $(\mu_n)_n$ tending to $+\infty$. Let $(\phi_n^\dagger)_n$ be the sequence of the corresponding equilibrium thresholds, which is increasing by Step 1, and let $\phi_\infty^\dagger \in (\phi^*, 1]$ be its limit. From equation (31), we have that $\lim_{n \rightarrow \infty} V_n'(\phi) = 0$ for all $\phi \in (\phi^*, \phi_\infty^\dagger)$. Hence, $\lim_{n \rightarrow \infty} V_n(\phi) = h - \ell$ for all $\phi \in (\phi^*, \phi_\infty^\dagger)$.

²⁷ Indeed, if there is a solution V of (31) and a value of ϕ^\dagger such that equation (44) is satisfied, there is an equilibrium as in Proposition 1 (since the smooth pasting conditions hold at ϕ^\dagger for some continuation value from the right given in (34) for some C_1 and with $C_2 = 0$).

If $\phi_\infty^\dagger = 1$ then the result holds. If $\phi_\infty^\dagger < 1$ then, given that the drift of the belief process from rejecting offers at each $\phi \in (\phi_\infty^\dagger, 1)$ becomes arbitrarily large as $\mu \rightarrow \infty$, we have that $\lim_{n \rightarrow \infty} V_n(\phi) = h - \ell$ for all $\phi \in (\phi_\infty^\dagger, \phi^*)$, hence the result holds.

(ii) Assume that, instead, $\bar{\mu} < +\infty$ (that is, $\phi^* < 1/2$). Then, if μ is large enough, the equilibrium characterization in Proposition 2 applies. By the same argument as when $\bar{\mu} = +\infty$, we now have $\lim_{\mu \nearrow +\infty} V'(\phi) = 0$ for all $\phi \in (\phi^*, 2\phi^*)$.

For each $\phi \in (2\phi^*, \hat{\phi}^{\dagger\dagger})$, we have (from equation (8))

$$V'(\phi) = -\frac{rh/2}{\mu(1-\phi)\phi\alpha(\phi)} \geq -\frac{rh}{4\mu(1-\phi)\phi^*},$$

where we used $\alpha(\phi) \geq 2\phi^*/\phi$ by the optimality of the seller's strategy.

We can now take a sequence $(\mu_n)_n$ tending to $+\infty$ and let $(\hat{\phi}_n^{\dagger\dagger})_n$ be the sequence of the corresponding equilibrium thresholds. Taking a subsequence if necessary, assume that $(\hat{\phi}_n^{\dagger\dagger})_n$ tends to some value $\hat{\phi}_\infty^{\dagger\dagger} \in [\phi^*, 1]$. Then, we can then use the same argument as for the case where $\bar{\mu} = +\infty$.

Step 4. We now want to show that $\lim_{\mu \rightarrow 0} V(\phi) = 0$ for all $\phi \in (\phi^*, 1)$. When μ is small enough, we can use the equilibrium characterization in Proposition 1. Take a sequence $(\mu_n)_n$ tending to 0. Let $(\phi_n^\dagger)_n$ be the sequence of the corresponding equilibrium thresholds. Taking a subsequence if necessary, assume that $(\phi_n^\dagger)_n$ tends to some value $\phi_\infty^\dagger \in [\phi^*, 1]$.

Now, for all $\phi \in (\phi^*, \phi_\infty^\dagger)$, we have $\lim_{n \rightarrow \infty} V_n'(\phi) = +\infty$ (from equation (31)). This has the implication that, $\phi_\infty^\dagger = \phi^*$. Because, for each n , the buyer is willing to reject all offers for all $\phi > \phi_n^\dagger$ for all n , and because the learning speed tends to 0 for all ϕ , the result holds. \square

Proof of Lemma 2

Proof. That Lemma 3 holds in the case where offers are not observable can be proven analogously. Additionally, by the usual take-it-or-leave-it offer argument, each seller either offers ℓ (which is accepted for sure by both types of the buyer) or the price $p(\phi) > \ell$ given in equation (17) (which is accepted for sure by the h -buyer and rejected for sure by the ℓ -buyer). \square

Proof of Proposition 4

Proof. The result follows trivially from Lemma 2 when $\phi \leq \phi^* := \ell/h$. The rest of the proof is divided into six steps.

Step 1: Preliminary results on regimes. Using that $\bar{\alpha}_\ell(\phi) = 1 - \pi(\phi)$ and $\bar{\alpha}_h(\phi) = 1$, the continuation

value satisfies

$$V(\phi) = r(\pi(\phi)(h - p(\phi)) + (1 - \pi(\phi))(h - \ell)) - \pi(\phi)\mu(1 - \phi)^2 \phi V'(\phi) + \frac{1}{2} \pi(\phi)\mu(1 - \phi)^2 \phi^2 V''(\phi). \quad (45)$$

Additionally, the h -buyer's indifference between accepting $p(\phi)$ or not implies that equation (17) holds as well.

We begin with a result characterizing some of the characteristics of the different possible regimes. The following result is analogous to Lemma 5 in the proofs of Propositions 1 and 2.

Lemma 6. *Assume prices are not observable. The following statements follow for all $\phi^* \leq \phi_1 < \phi_2 \leq 1$:*

(i) *If $\pi(\phi) \in (0, 1)$ for all $\phi \in (\phi_1, \phi_2)$ then $\phi_2 < 1$, $p(\phi) = \ell/\phi$, and*

$$\pi(\phi) = -\frac{r(h - \ell/\phi)}{(1 - \phi)\phi\mu V'(\phi)}. \quad (46)$$

(ii) *If $\pi(\phi) = 1$ for all $\phi \in (\phi_1, \phi_2)$ then*

$$p(\phi) = h + \frac{\mu}{r}(1 - \phi)\phi V'(\phi). \quad (47)$$

Proof. (i) The indifference of the sellers implies that $p(\phi) = \ell/\phi$. As a result, from equations (17) and (45), we have

$$0 = r(\phi - \phi^*)\frac{V'(\phi)}{h} + \frac{r}{2}(\phi - \phi^*)^2\frac{V''(\phi)}{h} + \phi^2\mu\left(1 - \phi^* - \frac{V(\phi)}{h}\right)\frac{V'(\phi)^2}{h^2} \quad (48)$$

holds. Equation (46) follows from equations (45) (with $p(\phi) = \ell/\phi$) and (48). Note that equation (46) can be written using log-likelihoods as

$$\pi(z) = -\frac{r(h - (1 + e^{-z})\ell)}{\mu V'(z)}.$$

Since it must be that $V'(z) \rightarrow 0$ as $z \rightarrow \infty$ (because V is bounded below by 0), we have that $\phi_2 < 1$.

(ii) **Case $\pi(\phi) = 1$ for all $\phi \in (\phi_1, \phi_2)$:** Equation (47) follows from the indifference of the h -buyer on accepting $p(\phi)$ or not, that is, equation (17). Additionally, we have the equation (45) becomes

$$rV(\phi) = \mu(1 - \phi)^2 \phi V'(\phi) + \frac{1}{2} \mu(1 - \phi)^2 \phi^2 V''(\phi). \quad (49)$$

□

(End of the proof of Lemma 6. Proof of Proposition 4 continues.)

Step 2: Continuity of p and π . In this step, we prove that p and π are continuous. Note that, since equation (17) holds for all ϕ , the continuity of π implies the continuity of p . Hence, we prove that π is continuous by contradiction. There are then two cases:

- (i) Assume first that $\pi(\phi) \in (0, 1)$ for some $\phi \in (0, 1)$ and there is a sequence $(\phi_n)_n$ converging to ϕ such that $\pi(\phi_n) = 1$ for all k . Note that $\pi(\phi) < 1$ only if $p(\phi) = \ell/\phi$. Then, we have that

$$\lim_{n \rightarrow \infty} p(\phi_n) = h + \frac{\mu}{r}(1-\phi)\phi V'(\phi) = \frac{\ell}{\phi} - \frac{1-\pi(\phi)}{\pi(\phi)}(h-\ell/\phi) < \ell/\phi.$$

This is a contradiction.

- (ii) The alternative case is that $\pi(\phi) \in [0, 1]$ and there is a sequence $(\phi_n)_n$ converging to ϕ such that $\pi(\phi_n) < 1$ for all k and $(\pi(\phi_n))_n$ converges to some $\bar{\pi} \neq \pi(\phi)$. Note that $p(\phi_n) = \ell/\phi_n$ for all n . Then, we have that

$$\ell/\phi = \lim_{n \rightarrow \infty} p(\phi_n) = h + \frac{\mu}{r}(1-\phi)\phi \bar{\pi} V'(\phi) = h - (h-p(\phi)) \frac{\bar{\pi}}{\pi(\phi)}.$$

If $\pi(\phi) = 1$ then, since $p(\phi) \geq \ell/\phi$, the right-hand side of the rightmost equality is strictly bigger than ℓ/ϕ , a contradiction. If instead $\pi(\phi) < 1$ then $p(\phi) = \ell/\phi$, but then the previous equation implies $\pi(\phi) = \bar{\pi}$, again a contradiction.

Step 3: Determination of the equilibrium structure. Let (ϕ_1, ϕ_2) be a maximal region with $\pi(\phi) \in (0, 1)$ for all $\phi \in (\phi_1, \phi_2)$ (that is, $\pi(\phi_1) \in \{0, 1\}$ and $\pi(\phi_2) = 1$). We aim at proving that there is at most one such interval, and is such that $\phi_1 = \phi^*$. Note that $\phi_2 < 1$ by Lemma 6. We have that $p(\phi_2) = \ell/\phi_2$. Then, using $\pi(\phi_2) = 1$, we have

$$\lim_{\phi \nearrow \phi_2} \pi'(\phi_2) = \frac{2\ell + \phi_2(2V(\phi_2) - h - \ell)}{h\phi_2(1-\phi_2)(\phi_2 - \phi^*)}. \quad (50)$$

We first argue that there is a unique value $\bar{\phi}_2$ such that the right-hand side of the previous equation is positive if $\phi_2 > \bar{\phi}_2$ and negative if $\phi_2 < \bar{\phi}_2$. To see this, recall V is decreasing. Hence, if the numerator is 0 for some value $\bar{\phi}_2$, it must be that $2V(\bar{\phi}_2) - h - \ell < 0$. If $\phi_2 > \bar{\phi}_2$, the numerator of the right-hand side of (50) is negative, while if $\phi_2 < \bar{\phi}_2$ the numerator is positive. Now, note that since $\pi'(\phi_2) \geq 0$ and $\phi_2 < 1$, we have that $\phi_2 \leq \bar{\phi}_2$. Also, if $\pi(\phi_1) = 1$, we have $\phi_1 \geq \bar{\phi}_2$ since $\pi'(\phi_1) \leq 0$. It then follows that there is at most one interval (ϕ_1, ϕ_2) where $\pi(\phi) \in (0, 1)$ for all $\phi \in (\phi_1, \phi_2)$ and $\pi(\phi_1), \pi(\phi_2) \in \{0, 1\}$. Furthermore, if such an interval exists, we have that $\phi_1 = \phi^*$ (and so $\pi(\phi_1) = 0$) and $\pi(\phi_2) = 1$.

We now prove that there must be some $\hat{\phi}^\ddagger \in (\phi^*, 1)$ such that $\pi(\phi) \in (0, 1)$ for all $\phi \in (\phi^*, \hat{\phi}^\ddagger)$. To show this, assume not; that is, assume that $\pi(\phi) = 1$ for all $\phi \in (\phi^*, 1)$. In this case, equation (49) holds for all $\phi \in (\phi^*, 1)$. The general solution to this equation is (34) for some $C_1, C_2 \in \mathbb{R}$ and where $\kappa > 0$ is defined in equation (35). Since V is bounded, it must be that $C_2 = 0$. The value-matching condition imposes that $V(\phi^*) = h - \ell$; that is,

$$V(\phi) = \left(\frac{1-\phi}{\phi} / \frac{1-\phi^*}{\phi^*} \right)^\kappa (h - \ell).$$

This implies that

$$\lim_{\phi \searrow \phi^*} p(\phi) = h - \frac{\mu}{r} \kappa (h - \ell) < h.$$

This is a contradiction, since $p(\phi) \geq \ell / \phi$.

We let $\hat{\phi}^\ddagger \in (\phi^*, 1)$ be the supremum value such that $\pi(\phi) \in (0, 1)$ for all $\phi \in (\phi^*, \hat{\phi}^\ddagger)$.

Step 4. Equations for existence and uniqueness. From Step 3, we have that there is some $C_1 \in \mathbb{R}_{++}$ such that, for all $\phi \in (\phi^*, 1)$, we have²⁸

$$V(\phi; C_1) = C_1 \left(\frac{1-\phi}{\phi} \right)^\kappa. \quad (51)$$

For each C_1 , we let $\hat{\phi}(C_1)$ be determined by setting $\pi(\hat{\phi}(C_1)) = 1$ in equation (46) replacing $V'(\phi)$ by $V'(\hat{\phi}(C_1); C_1)$. We then have

$$C_1 = \frac{r \left(\frac{1-\hat{\phi}(C_1)}{\hat{\phi}(C_1)} \right)^{-\kappa} (\hat{\phi}(C_1) - \phi^*) h}{\mu \kappa \hat{\phi}(C_1)}. \quad (52)$$

The derivative of the right-hand side of the previous expression with respect to $\hat{\phi}(C_1)$ is

$$\frac{r \left(\frac{1-\hat{\phi}(C_1)}{\hat{\phi}(C_1)} \right)^{-\kappa} (\phi^* (1-\phi^*) - \overbrace{(\phi^* - \kappa)(\hat{\phi}(C_1) - \phi^*)}^{(*)}) h}{\mu \kappa (1-\hat{\phi}(C_1)) \hat{\phi}(C_1)^2}.$$

Given that $\hat{\phi}(C_1) > \phi^*$ and $\kappa > 0$, the term $(*)$ is bounded above by $\phi^* (1-\phi^*)$. Then, for each $C_1 > 0$, $\hat{\phi}(C_1)$ is uniquely defined, and $\hat{\phi} : \mathbb{R}_{++} \rightarrow (\phi^*, 1)$ is a bijective function. Hence, from now on, abusing notation, we use $\hat{\phi} \in (\phi^*, 1)$ instead of C_1 as the free variable coming from the differential equation (49). We also use $V(\cdot; \hat{\phi})$ instead of $V(\cdot; C_1)$.

²⁸ Note that this coincides with expression (34) with $C_2 = 0$, satisfied by the continuation value on the upper belief regions of beliefs in the observable case.

For any given $\hat{\phi}$, the pair of values of $V(\hat{\phi}; \hat{\phi})$ and $V'(\hat{\phi}; \hat{\phi})$ can be used as boundary conditions to obtain a unique solution to equation (48) on $(\phi^*, \hat{\phi}]$, which we denote $V(\phi; \hat{\phi})$ without risk of confusion.²⁹ From the previous expressions we have that

$$V(\hat{\phi}; \hat{\phi}) = \frac{(\hat{\phi} - \phi^*) r h}{\hat{\phi} \kappa \mu}. \quad (53)$$

The right hand side of equation (53) is increasing in $\hat{\phi}$ on $(\phi^*, 1)$, and 0 at $\hat{\phi} = \phi^*$. Additionally, the previous expressions imply

$$V'(\hat{\phi}; \hat{\phi}) = -\frac{(\hat{\phi} - \phi^*) r h}{(1 - \hat{\phi}) \hat{\phi}^2 \mu}. \quad (54)$$

It is easy to see that this is decreasing in $\hat{\phi}$ for all $\hat{\phi} \in (\phi^*, 1)$.

Hence, there is an equilibrium (which must be as specified in the statement of the proposition) for a given value $\phi^\ddagger \in (\phi^*, 1)$ only if there is a solution of equation (48) on (ϕ^*, ϕ^\ddagger) , denoted $V(\cdot; \phi^\ddagger)$, with boundary conditions (53) and (54), and the lower boundary condition holds, that is, if $\lim_{\phi \searrow \phi^*} V(\phi; \phi^\ddagger) = h - \ell$.

Step 5. Change of variables. From now on, we assume an equilibrium with continuation value V exists, and we will establish the necessary and sufficient conditions that it satisfies, and we will finally establish the existence of a unique equilibrium in Step 6. We change variables defining, for each $\hat{\phi} \in (\phi^*, 1)$,

$$W(y) := \frac{2^{1/2} \sqrt{\mu/r} \hat{\phi}}{(\hat{\phi} - \phi^*) h} \left(V\left(\frac{\hat{\phi} \phi^* y}{\hat{\phi}(y-1) + \phi^*}\right) - (h - \ell) \right)$$

for all $y \in [1, +\infty)$. Note that W is negative and increasing. Note also that the limit $\phi \searrow \phi^*$ corresponds to the limit $y \rightarrow \infty$, while $\phi = \hat{\phi}$ corresponds to $y = 1$. Using our definition, equation (48) takes the following simpler form:

$$W''(y) = y^2 W(y) W'(y)^2. \quad (55)$$

Since $W(\cdot)$ is negative, it is also concave. The requirement that $\lim_{\phi \searrow \phi^*} V(\phi) = h - \ell$ corresponds to $\lim_{y \rightarrow \infty} W(y) = 0$. We will now analyze solutions to equation (55).

We note that the boundary conditions (53) and (54) can be written as boundary conditions on W at $y = 1$ as follows:

²⁹ Note that, by the smooth pasting condition, $V(\cdot; \hat{\phi})$ is continuous and differentiable at $\hat{\phi}$. It satisfies equation (48) on $(\phi^*, \hat{\phi}]$ and equation (49) on $[\hat{\phi}, 1)$.

$$W(1; \hat{\phi}) = 2^{1/2} \sqrt{\mu/r} \left(\frac{r}{\kappa\mu} - \frac{\hat{\phi}(1-\phi^*)}{\hat{\phi}-\phi^*} \right), \quad (56)$$

$$W'(1; \hat{\phi}) = \frac{2^{1/2}}{\sqrt{\mu/r}} \frac{\hat{\phi}-\phi^*}{(1-\hat{\phi})\phi^*}. \quad (57)$$

The right-hand side of equation (56) is strictly negative if and only if $\hat{\phi} \in (\phi^*, \hat{\phi}^+)$, where

$$\hat{\phi}^+ := \min \{ 1, \phi^* (1 - (1 - \phi^*) \kappa \mu / r)^{-1} \} > \phi^*.$$

Also, in this range, $W(1; \hat{\phi})$ increases from $-\infty$ to either 0 (if $\hat{\phi}^+ < 1$) or $2^{1/2}(\kappa^{-1} - \sqrt{\mu/r})$ (if $\hat{\phi}^+ = 1$). The right-hand side of equation (57) is strictly positive and strictly increasing in $\hat{\phi}$ on $(\phi^*, 1)$, and increases from 0 to $+\infty$.

Step 6. Existence and uniqueness. We now aim to show that there is a unique $\hat{\phi} \in (\phi^*, 1)$ such that

$$\lim_{y \nearrow \infty} W(y; \hat{\phi}) = 0,$$

where $W(\cdot; \hat{\phi})$ is the solution to (55) with boundary conditions given by equations (56) and (57). We do it in two parts:

- (i) **Uniqueness:** Take two different values $\hat{\phi}_1$ and $\hat{\phi}_2$ satisfying $\phi^* < \hat{\phi}_1 < \hat{\phi}_2 < \hat{\phi}^+$. By the previous results, $W(1; \hat{\phi}_1) < W(1; \hat{\phi}_2) < 0$ and $0 < W'(1; \hat{\phi}_1) < W'(1; \hat{\phi}_2)$. We want to show that $W(y; \hat{\phi}_2) - W(y; \hat{\phi}_1)$ is increasing in y , and so $W(\cdot; \hat{\phi})$ tends to 0 for at most one value of $\hat{\phi} \in (\phi^*, \hat{\phi}^+)$. Assume, for the sake of contradiction, there is some value y' such that $W'(y'; \hat{\phi}_1) = W'(y'; \hat{\phi}_2)$, and let y be the infimum with this property. It then has to be that $W(y; \hat{\phi}_1) < W(y; \hat{\phi}_2)$. Then we have

$$W''(y; \hat{\phi}_1) = y^2 W(y; \hat{\phi}_1) W'(y; \hat{\phi}_1)^2 < y^2 W(y; \hat{\phi}_2) W'(y; \hat{\phi}_2)^2 = W''(y; \hat{\phi}_2).$$

This is a contradiction, since $W'(\cdot; \hat{\phi}_1) < W'(\cdot; \hat{\phi}_2)$ on $(1, y)$. Hence, we have that $W'(y; \hat{\phi}_1) < W'(y; \hat{\phi}_2)$ for all $y > 1$. Therefore, if $W(y; \hat{\phi}_1)$ and $W(y; \hat{\phi}_2)$ are convergent as $y \rightarrow \infty$, they converge to different values. A similar argument implies that two solutions of equation (55) cross at most once.

- (ii) **Existence:** Note that a particular solution of (55) is $\hat{W}(y) := -2^{1/2}/y$.³⁰ Note also that, as $\hat{\phi} \rightarrow \phi^*$,

$$W(1; \hat{\phi}) \rightarrow -\infty < \hat{W}(1) \quad \text{and} \quad W'(1; \hat{\phi}) \rightarrow 0 < \hat{W}'(1).$$

Hence, if $\hat{\phi}$ is close enough to ϕ^* , $\lim_{y \rightarrow \infty} W(y; \hat{\phi})$ is strictly lower than $\lim_{y \rightarrow \infty} \hat{W}(y)$, which is

³⁰ Such solution is an equilibrium when $\mu = r$, in which case $\hat{\phi}^\dagger = 2\phi^*/(1 + \phi^*)$.

equal to 0. We assume, for the sake of contradiction, that $\lim_{y \rightarrow \infty} W(y; \hat{\phi}^+) = w^+$, for some $w^+ < 0$. Since $W(1; \hat{\phi}^+) = 0$ if $\hat{\phi}^+ < 1$, then it must be that $\hat{\phi}^+ = 1$ and so $W'(1; \hat{\phi}^+) = +\infty$. Then, any solution of equation (55) with $W(1) < W(1; \hat{\phi}^+)$ is such that $\lim_{y \rightarrow \infty} W(y) < w^+$. We let $W(y; 1)$ be defined as $\lim_{\hat{\phi} \rightarrow 1} W(y; \hat{\phi})$, which by assumption satisfies $\lim_{y \rightarrow \infty} W(y; 1) < 0$. For each ε , let $\tilde{W}_\varepsilon(\cdot)$ be defined as the solution to equation (55) satisfying $\tilde{W}_\varepsilon(2) = W(2; 1)$ and $\tilde{W}'_\varepsilon(2) = W'(2; 1) + \varepsilon$. It is clear that if $\varepsilon > 0$ is chosen strictly larger than 0, then $\tilde{W}_\varepsilon(1) < W(1, 1)$. Since solutions to (55) only cross once, we have $\lim_{y \rightarrow \infty} \tilde{W}_\varepsilon(y) > w^+$, but this is a contradiction. Hence, there exists a unique value $\hat{\phi}^+$ such that $W(y; \hat{\phi}^+) = 0$, and hence a unique equilibrium exists. \square

Proof of Proposition 5

Proof. Recall that, as we proved in the proof of Proposition 4, the equilibrium is divided into three regions. In the region $(0, \phi^*)$, $V = h - \ell$. In the region (ϕ^*, ϕ^\ddagger) , V follows equation (48) (with $\lim_{\phi \searrow \phi^*} V(\phi) = h - \ell$) and π satisfies equation (46). In the region (ϕ^*, ϕ^\ddagger) , V satisfies equation (34) with $C_2 = 0$, and $\pi(\phi) = 1$. Recall also that $\lim_{\phi \nearrow \phi^\ddagger} \pi(\phi) = 1$.

We first note that we can define

$$U(\phi) := \sqrt{\mu/r} \left(\frac{V(\phi)}{h} - 1 + \phi^* \right). \quad (58)$$

so that equation (48) becomes

$$U''(b) = - \frac{\phi - \phi^* + \phi^2 U(\phi) U'(\phi)}{(\phi - \phi^*)^2},$$

with the condition now that $U(\phi^*) = 0$. Note that U is negative and does not depend on μ . Note now that we can rewrite equation (46) as

$$\pi(\phi) = \frac{\phi - \phi^*}{\sqrt{\mu/r} (1 - \phi) \phi^2 (-U'(\phi))}.$$

Since π is increasing, $\pi(\phi^\ddagger) = 1$, and the right-hand side of the previous equation is decreasing in μ , it follows that ϕ^\ddagger is increasing in μ (recall Figure 3(b)).

We define

$$\hat{V}_\mu(\phi) := h - \ell + \frac{1}{\sqrt{\mu/r}} h U(\phi). \quad (59)$$

(Note that $V(\phi) = \hat{V}_\mu(\phi)$ for $\phi \in (\phi^*, \phi^\ddagger]$. Note that $\hat{V}_\mu(\phi)$ is increasing in μ (since $U(\phi)$ is negative).

Define also

$$\tilde{V}_\mu(\phi; C_\mu) := C_\mu \left(\frac{1-\phi}{\phi} \right)^{\kappa(\mu)},$$

where $\kappa(\mu)$ is the right-hand side of (35). (Note that $V(\phi) = \tilde{V}_\mu(\phi; C_\mu)$ when $\phi \in [\phi^\ddagger, 1)$, where C_μ is the value of C_1 in equation (34) for the unique equilibrium.) By letting ϕ_μ^\ddagger be the value of ϕ^\ddagger for μ , note also that the smooth pasting condition implies $\hat{V}_\mu(\phi_\mu^\ddagger) = \tilde{V}_\mu(\phi_\mu^\ddagger; C_\mu)$ and $\hat{V}'_\mu(\phi_\mu^\ddagger) = \tilde{V}'_\mu(\phi_\mu^\ddagger; C_\mu)$.

Take $\mu_1, \mu_2 \in \mathbb{R}_{++}$ with $\mu_1 < \mu_2$. Using subindexes to denote the variables of the corresponding (unique) equilibria, the previous observations imply that $\phi_{\mu_1}^\ddagger < \phi_{\mu_2}^\ddagger$ and that $V_{\mu_1}(\phi) < V_{\mu_2}(\phi)$ for all $\phi \in (\phi^*, \phi_{\mu_1}^\ddagger]$. Assume, for the sake of contradiction, that there is some $\hat{\phi} \in (\phi_{\mu_1}^\ddagger, 1)$ such that $V_{\mu_1}(\hat{\phi}) = V_{\mu_2}(\hat{\phi})$. There are two cases:

- (i) Consider first the case $\hat{\phi} \in (\phi_{\mu_1}^\ddagger, \phi_{\mu_2}^\ddagger)$. We then have $\tilde{V}_{\mu_1}(\hat{\phi}, C_{\mu_1}) = \hat{V}_{\mu_2}(\hat{\phi})$ and $\tilde{V}'_{\mu_1}(\hat{\phi}, C_{\mu_1}) \geq \hat{V}'_{\mu_2}(\hat{\phi})$. Let \hat{C}_{μ_2, μ_1} be the value such that $\tilde{V}_{\mu_2}(\hat{\phi}, \hat{C}_{\mu_2, \mu_1}) = \tilde{V}_{\mu_1}(\hat{\phi}, C_{\mu_1})$, that is,

$$\hat{C}_{\mu_2, \mu_1} = C_\mu \left(\frac{1-\hat{\phi}}{\hat{\phi}} \right)^{\kappa(\mu_1) - \kappa(\mu_2)}.$$

Simple algebra shows that

$$\tilde{V}'_{\mu_2}(\hat{\phi}, \hat{C}_{\mu_2, \mu_1}) = \frac{\kappa(\mu_2)}{\kappa(\mu_1)} \tilde{V}'_{\mu_1}(\hat{\phi}, C_{\mu_1}).$$

Since $\kappa(\cdot)$ is a decreasing function, we have that $\tilde{V}'_{\mu_2}(\hat{\phi}, \hat{C}_{\mu_2, \mu_1})$ is smaller in absolute value than $\tilde{V}'_{\mu_1}(\hat{\phi}, C_{\mu_1})$, hence it is higher because both are negative. Therefore, using that $\tilde{V}'_{\mu_1}(\hat{\phi}, C_{\mu_1}) \geq \hat{V}'_{\mu_2}(\hat{\phi})$, we have $\tilde{V}'_{\mu_2}(\hat{\phi}, \hat{C}_{\mu_2, \mu_1}) > \hat{V}'_{\mu_2}(\hat{\phi})$. It then follows that there exists some $\hat{C}'_{\mu_2} < \hat{C}_{\mu_2, \mu_1}$ and $\hat{\phi}' < \hat{\phi}$ such that $\hat{V}_{\mu_2}(\hat{\phi}') = \tilde{V}_{\mu_2}(\hat{\phi}', \hat{C}'_{\mu_2})$ and $\hat{V}'_{\mu_2}(\hat{\phi}') = \tilde{V}'_{\mu_2}(\hat{\phi}', \hat{C}'_{\mu_2})$. Nevertheless, this implies that there exists an equilibrium for μ_2 with $\phi^\ddagger = \hat{\phi}' < \phi_{\mu_1}^\ddagger$, which contradicts the uniqueness of the equilibrium established in Proposition 4.

- (ii) Consider now the case $\hat{\phi} \in [\phi_{\mu_2}^\ddagger, 1)$. In this case we have $\tilde{V}_{\mu_1}(\hat{\phi}; C_{\mu_1}) = \tilde{V}'_{\mu_2}(\hat{\phi}, C_{\mu_2})$ and $\tilde{V}'_{\mu_1}(\hat{\phi}; C_{\mu_1}) \geq \tilde{V}'_{\mu_2}(\hat{\phi}, C_{\mu_2})$. This implies that $C_{\mu_2} = \hat{C}_{\mu_2, \mu_1}$ defined above. As we argued, we have that $\tilde{V}_{\mu_1}(\hat{\phi}; C_{\mu_1}) < \tilde{V}'_{\mu_2}(\hat{\phi}, C_{\mu_2})$, which is a contradiction. \square

Proof of Proposition 6

Proof. When $\mu \geq \bar{\mu}$ the result is trivial, so we focus on the case $\mu < \bar{\mu}$. We define

$$\hat{V}^{\text{ob}}(\phi) = h - \ell + \frac{r(\phi h - \ell)}{\mu \phi} + \frac{r(h - \ell)}{\mu} \log\left(\frac{1-\phi}{\phi} / \frac{1-\phi^*}{\phi^*}\right)$$

for all $\phi \in [\phi^*, 1)$. Note that $\hat{V}^{\text{ob}}(\phi) = V^{\text{ob}}(\phi)$ for all $\phi \in [\phi^*, \phi^\dagger]$ (recall equation (31)), but $\hat{V}^{\text{ob}}(\phi) > V^{\text{ob}}(\phi)$ for all $\phi \in (\phi^\dagger, 1]$. As the proof of Propositions 1 and 2 argues, $\frac{\hat{V}^{\text{ob}}(\phi)}{-\hat{V}^{\text{ob}'(\phi)}} > \kappa^{-1}$ for all $\phi < \phi^\dagger$ and $\frac{\hat{V}^{\text{ob}}(\phi^\ddagger)}{-\hat{V}^{\text{ob}'(\phi^\ddagger)}} < \kappa^{-1}$ for all $\phi > \phi^\dagger$.

We also define $z^\dagger = \check{z}(\phi^\dagger)$ and $z^\ddagger = \check{z}(\phi^\ddagger)$ (recall the definition of \check{z} in (20)). Recall that equation (44) holds for V^{ob} , and we also have

$$\frac{V^{\text{un}}(z)}{-V^{\text{un}'(z)}} \begin{cases} > \kappa^{-1} & \text{if } z \in (z^*, z^\ddagger), \\ = \kappa^{-1} & \text{if } z \in [z^\ddagger, +\infty). \end{cases} \quad (60)$$

Note that $\hat{V}^{\text{ob}'(\phi^\ddagger)} = V^{\text{un}'(\phi^\ddagger)}$ (because $\pi^{\text{un}}(\phi^\ddagger) = 1$, and so $\gamma^{\text{ob}}(\phi^\ddagger) = \gamma^{\text{un}}(\phi^\ddagger) = 1$ in equation (18)), and so $\hat{V}^{\text{ob}}(\phi^\ddagger) > V^{\text{un}}(\phi^\ddagger)$. As a result, $\frac{\hat{V}^{\text{ob}}(\phi^\ddagger)}{-\hat{V}^{\text{ob}'(\phi^\ddagger)}} > \kappa^{-1}$, hence it must be that $\phi^\ddagger < \phi^\dagger$.

The argument proceeds as in the proof of Proposition 3. Assume, for the sake of contradiction, that there is some $z > z^\ddagger$ such that $V^{\text{un}}(z) = V^{\text{ob}}(z)$. Since $V^{\text{un}}(z^\ddagger) < V^{\text{ob}}(z^\ddagger)$, there must then be some $\hat{z} > z^\ddagger$ such that $V^{\text{un}}(\hat{z}) = V^{\text{ob}}(\hat{z})$ and $-V^{\text{un}'(\hat{z})} \leq -V^{\text{ob}'(\hat{z})}$. But then, this implies,

$$\frac{V^{\text{ob}}(\hat{z})}{-V^{\text{ob}'(\hat{z})}} \leq \frac{V^{\text{un}}(\hat{z})}{-V^{\text{un}'(\hat{z})}} = \kappa^{-1},$$

that is, $\frac{V^{\text{ob}}(\hat{z})}{-V^{\text{ob}'(\hat{z})}} = \kappa^{-1}$. This implies that $\hat{z} \geq z^\dagger$. Nevertheless, we then have that $V^{\text{un}}(z^\dagger) < V^{\text{ob}}(z^\dagger)$. Since V^{un} and $V^{\text{ob}}(z)$ follow the same equation for $z > z^\dagger$, this implies that $V^{\text{un}}(z) < V^{\text{ob}}(z)$ for all $z > z^\dagger$, which is a contradiction. This concludes the proof of the proposition. \square

Proof of Proposition 7

Proof. We prove each part separately:

- (i) We first prove that, if $\mu \leq \bar{\mu}$, then $W^{\text{ob}}(\phi_0) > W^{\text{un}}(\phi_0) > W^{\text{no}}(\phi_0)$ for all $\phi_0 \in (\phi^*, 1)$. The last inequality is trivial for the reasons laid out in the main text after the proposition. The first inequality is obtained as follows. Note that, when $\mu \leq \bar{\mu}$, we have that $\alpha^x(\phi) = 1$ for all $x \in \{\text{un}, \text{ob}\}$. Hence, the term $(**)$ in equation (19) is equal to h for all x . As explained after the proposition, the term $(*)$ in equation (19) is equal to $\frac{\ell}{h-\ell} V^x(\phi_0)$ for all $x \in \{\text{un}, \text{ob}\}$. Then, applying Proposition 6, the result follows.
- (ii) We now prove that, if $\mu > \bar{\mu}$, then $W^{\text{un}}(\phi_0) > W^{\text{no}}(\phi_0)$ for all $\phi_0 \in (\phi^*, 1)$. This result holds trivially by the arguments after the proposition. \square

B Discrete-time approximation

This Appendix contains some derivations of discrete-time versions of the model in Section 2. We illustrate how outcomes can be obtained as the limit of outcomes of discrete-time models as the length of the period length shrinks to 0.

We begin presenting a continuous-time analogous to our model in Section 2 where the buyer meets a seller only at times belonging to a discrete set. Formally, at each time t in the discrete set $\mathcal{T}_\Delta := \{0, \Delta, 2\Delta, \dots\} \subset \mathbb{R}_+$, for some small $\Delta > 0$, the buyer meets a short-lived seller, the “ t -seller”. First, the t -seller offers price p_t , and then the buyer decides to purchase or not, denoted $a_t = 1$ and $a_t = 0$, respectively. The buyer values the goods of all sellers the same. The valuation (or type) is private, and it is either ℓ or h with $0 < \ell < h$. The buyer discounts the future at a discount factor $e^{-\Delta r}$. For a given history of acceptance decisions and prices $(a_t, p_t)_{t \in \mathcal{T}_\Delta}$, the payoff of the θ -buyer is

$$\sum_{t \in \mathcal{T}_\Delta} e^{-tr} (1 - e^{-\Delta r}) a_t (\theta - p_t).$$

The payoff of the t -seller is $a_t p_t$.

There is a public signal about the previous purchasing decisions of the buyer. The t -seller observes a history of signals $(x_{t'})_{t' \in \mathcal{T}_\Delta, t' < t}$. The value of x_t depends on the acceptance decision of the buyer at time t as follows according to a normal distribution:

$$x_t \sim N(\Delta \mu a_t, \Delta^{1/2}).$$

Hence, each x_t is informative of the buyer’s acceptance decision a_t . Higher realizations of x_t are indicative of acceptance. Note the previous specification is equivalent to assuming that t -seller observes $(X_{t'})_{t' \in \mathcal{T}_\Delta, t' < t}$, where

$$X_t = \mu \sum_{t' \in \mathcal{T}_\Delta, t' < t} \Delta a_{t'} + B_t \tag{61}$$

and B_t is a normalized Wiener process.³¹ Such equivalence follows by defining $x_{t+\Delta} := X_{t+\Delta} - X_t$. Equation (61) is a discrete-time version of equation (1).³²

³¹ Recall that the total variation contains the information that can be obtained about a (constant) drift of a Brownian motion over an interval of time.

³² Note that the Bayes’ rule specifies that

$$\frac{\phi_{t+\Delta}}{1 - \phi_{t+\Delta}} = \frac{\phi_t}{1 - \phi_t} \frac{\alpha_h(\phi_t) f_1(x_t) + (1 - \alpha_h(\phi_t)) f_0(x_t)}{\alpha_\ell(\phi_t) f_1(x_t) + (1 - \alpha_\ell(\phi_t)) f_0(x_t)},$$

Fix a consistent strategy $(\alpha_\ell, \alpha_h, \tilde{\pi})$ (as defined in Section 2). Then, for any strategy $\hat{\alpha}$ consistent with $(\alpha_\ell, \alpha_h, \tilde{\pi})$ we let $V_\Delta(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi})$ and $V(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi})$ denote the implied continuation value of the h -buyer in the discrete with parameter Δ and the continuous model, respectively.

The following result establishes that the payoff the h -buyer obtains from a given strategy profile in the continuous model approximately coincides with the payoff he obtains from such strategy profile in the discrete-time model when Δ is small enough.

Proposition 8. *Let $(\alpha_\ell, \alpha_h, \tilde{\pi})$ be a piecewise-Lipschitz strategy profile and $\hat{\alpha}$ be a piecewise-Lipschitz buyer's strategy. Then, for all $\phi \in [0, 1]$, we have that*

$$\lim_{\Delta \searrow 0} V_\Delta(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi}) = V(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi}).$$

Proof. We fix some strategy profile $(\alpha_\ell, \alpha_h, \tilde{\pi})$ and a buyer strategy $\hat{\alpha}$. Under such strategy profile and strategy, the payoff of the h -buyer at time t —for simplicity denoted $V_\Delta(\phi) := V_\Delta(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi})$ —satisfies the following Bellman equation:

$$\begin{aligned} V_\Delta(\phi_t) = & \mathbb{E}_{\tilde{p} \sim \tilde{\pi}} \left[\hat{\alpha}(\phi_t, \tilde{p}) \left((1 - e^{-r\Delta})(h - \tilde{p}) + e^{-r\Delta} \mathbb{E}_{\tilde{\phi}} [V_\Delta(\tilde{\phi}) | \text{acc}, \tilde{p}, \phi_t] \right) \right. \\ & \left. + (1 - \hat{\alpha}(\phi_t, \tilde{p})) e^{-r\Delta} \mathbb{E}_{\tilde{\phi}} [V_\Delta(\tilde{\phi}) | \text{rej}, \tilde{p}, \phi_t] \right] | \phi_t \end{aligned} \quad (62)$$

where $\mathbb{E}_{\tilde{p} \sim \tilde{\pi}}[\cdot | \phi_t]$ denotes the expectation with respect to \tilde{p} (which is distributed according to $\tilde{\pi}(\phi_t)$), and where $\mathbb{E}_{\tilde{\phi}}[\cdot | \text{acc}, \tilde{p}, \phi_t]$ and $\mathbb{E}_{\tilde{\phi}}[\cdot | \text{rej}, \tilde{p}, \phi_t]$ denotes the expectation of next period's belief conditional on price \tilde{p} being offered and the buyer accepting and rejecting the offer, respectively (these distributions are obtained using standard Bayes rule). Standard fixed-point arguments imply that there is a unique $V_\Delta(\cdot)$ satisfying (62). Note that the previous equation can be rewritten as

$$V_\Delta(\phi_t) = (1 - e^{-r\Delta}) \mathbb{E}_{\tilde{p} \sim \tilde{\pi}} \left[\hat{\alpha}(\phi_t, \tilde{p})(h - \tilde{p}) \right] + e^{-r\Delta} \mathbb{E}_{\tilde{\phi}} [V(\tilde{\phi}) | \phi_t, \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi}], \quad (63)$$

where $\mathbb{E}_{\tilde{\phi}}[\cdot | \phi_t, \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi}]$ is the expectation using the distribution of $\phi_{t+\Delta}$ induced by $(\hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi})$ given ϕ_t .

where $f_m(\cdot)$ is the density of $N(\Delta \mu m, \Delta^{1/2})$, for all $m \in \mathbb{R}$. If, as in the main text, a_t is a number in $[0, 1]$ (interpreted as the mass of goods purchased by the consumer), then Bayes' rule is

$$\frac{\phi_{t+\Delta}}{1 - \phi_{t+\Delta}} = \frac{\phi_t}{1 - \phi_t} \frac{f_{a_h(\phi_t)}(x_t)}{f_{a_\ell(\phi_t)}(x_t)}.$$

It is easy to see that these two updating rules become equivalent as Δ shrinks to 0.

Simple algebra shows that the expected change in beliefs and its variance is given by

$$\begin{aligned}\mathbb{E}_{\tilde{\phi}}[\phi_{t+\Delta}-\phi_t|\phi_t,\hat{\alpha},\alpha_\ell,\alpha_h,\hat{p}]&=\tilde{\mu}(\phi_t,\hat{p};\hat{\alpha},\alpha_\ell,\alpha_h)\Delta+O(\Delta^2) \text{ and} \\ \mathbb{V}_{\tilde{\phi}}[\phi_{t+\Delta}-\phi_t|\phi_t,\hat{\alpha},\alpha_\ell,\alpha_h,\hat{p}]&=\tilde{\sigma}^2(\phi_t,\hat{p};\hat{\alpha},\alpha_\ell,\alpha_h)\Delta+O(\Delta^2),\end{aligned}$$

where $\tilde{\mu}$ and $\tilde{\sigma}^2$ are defined in equations (2) and (3), respectively. Standard arguments imply then that, as $\Delta \rightarrow 0$, ϕ_t approaches a stochastic process with drift and diffusion parameters given by (2) and (3), respectively. As a result, the first-order terms in Δ of equation (63) converge to equation (4). \square

An implication of Proposition 8 is the following. Let $(\alpha_\ell, \alpha_h, \tilde{\pi})$ be an equilibrium of the continuous model (as defined in Definition 1), which we showed to be unique in Propositions 1 and 2. From equation (6), it is clear that each seller behaves optimally given the buyer's behavior also in the discrete-time model. Using that $(\alpha_\ell, \alpha_h, \tilde{\pi})$ is an equilibrium and also using Proposition 8, it follows that, for each ϕ , each $\hat{\alpha} \neq \alpha_h$, and each $\varepsilon > 0$, we have $V_\Delta(\phi; \alpha_h, \alpha_\ell, \alpha_h, \tilde{\pi}) > V_\Delta(\phi; \hat{\alpha}, \alpha_\ell, \alpha_h, \tilde{\pi}) - \varepsilon$ if Δ is small enough. That is, if players play an equilibrium $(\alpha_\ell, \alpha_h, \tilde{\pi})$ of the continuous model, neither sellers nor the buyer can gain more than ε from deviating in a discrete-time version of the model if the length of the period is small enough.