Bargaining and News

Brendan Daley
Duke University
The Fuqua School of Business
bd28@duke.edu
http://faculty.fuqua.duke.edu/∼bd28

Brett Green
UC Berkeley
Haas School of Business
bgreen@haas.berkeley.edu
http://faculty.haas.berkeley.edu/bgreen/

April 18, 2016

Preliminary Draft

Abstract

We study a bargaining model in which a buyer makes frequent offers to a privately informed seller. In addition, the buyer gradually learns about the seller's type from a publicly observable news process. We show that the buyer's ability to leverage this information to extract more surplus from the seller is remarkably limited. In fact, the buyer's equilibrium payoff is identical to what he would achieve if he were unable to price discriminate. Delay occurs only if the adverse-selection problem is sufficiently strong. When trade is delayed, the buyer engages in a kind of costly "experimentation" by making offers that are sure to earn him negative payoffs if accepted, but that improve his information and expected continuation payoff if rejected. We investigate the effects of market power by comparing our results to a setting with competitive buyers. Surprisingly, both efficiency and the seller's payoff can decrease by introducing competition among buyers.

1 Introduction

A central issue in the bargaining literature is whether trade will be (inefficiently) delayed. What is often ignored, however, is that if trade is in fact delayed, new information may come to light.[1] Of course, the players' anticipation of this information may itself affect the amount of delay in their bargaining.

For example, consider a startup that has "catered" its innovation to a large firm with the aim of being acquired (an increasing common strategy in entrepreneurship—see Wang (2015)). The longer the startup operates as an independent business, the more the large firm expects to learn about the quality of the innovation, which can influence the offers that it tenders. At the same time, delay is inefficient as the large firm can generate greater value from the innovation due to economies of scale and possession of a portfolio of complementary products. We are interested in how the large firm's ability to learn about the startup over time affects its relative bargaining power, trading dynamics, and the amount of surplus realized from the potential acquisition.

In this paper, we study a model of bargaining in which the uninformed party (the "buyer") makes frequent offers to the informed party (the "seller") while simultaneously learning gradually about the seller's type from a publicly observable *news* process. There is common knowledge of gains from trade, values are (potentially) interdependent, and the seller is privately informed about the quality of the tradable asset (i.e., her type), which may be either high or low. Because of discounting, the uniquely efficient outcome is immediate trade. We pose the model directly in continuous time, which captures the idea that there are no institutional frictions in the bargaining protocol and facilitates a tractable analysis.[2] News is modeled as a Brownian diffusion process with type-dependent drift.

We show that the buyer's ability to leverage this information to extract more surplus from the seller is remarkably limited. In particular, the buyer's equilibrium payoff is identical to what he would achieve if he were unable to screen through prices (i.e., could only make offers that the seller surely finds acceptable or not, regardless of her type). Consequently, delay occurs if and only if the adverse-selection problem is sufficiently severe. This finding extends the no-delay result from the Coase conjecture literature (Coase, 1972; Fudenberg et al., 1985; Gul et al., 1986)—in which values are independent, meaning there is no adverse-selection problem—by showing that the Coasian incentive to speed up trade overwhelms the option value of waiting for gradual news release.

When trade is delayed due to adverse selection, the buyer engages in a form of costly "experimentation" by making offers that are sure to earn him negative payoffs if accepted, but that improve his information and expected continuation payoff if rejected. The buyer,

[1] Fuchs and Skrzypacz (2010) is a notable exception, as we will discuss.

[2] Ortner (2016) also poses a bargaining game directly in continuous time.

however, exhausts all of the benefits from this experimentation leaving him with precisely the same payoff he would obtain if he were unable to offer such prices. In fact, the main beneficiary of this experimentation is the low-type seller, who always earns more than her value to the buyer.

We investigate the effects of market power by comparing our results to those of the competitive-buyer model of Daley and Green (2012) (hereafter, DG12). We find novel differences in both the pattern of trade and the resulting efficiency. Perhaps surprisingly, both efficiency and the seller's payoff can *decrease* by introducing competition among buyers. These results stem from the fact that changes in the level of buyer competition alter the relative market power of buyers versus sellers, who have different expectations about the realization of future news.

Intuitively, the amount of delay is driven by the party that stands to gain from information release. With buyer competition, it is the high-value seller who gains from information release as buyers bid up the price after positive news realizations. Without this competition, the (single) buyer is the player who essentially determines the amount of delay. Because of her private information, the high-type seller is more optimistic about the realization of future news than is the buyer, which causes her to delay trade when facing competitive buyers in states where a single buyer would chose to trade immediately. This finding is illustrated starkly in the no-adverse-selection case: with a single buyer trade is immediate, whereas it will be delayed with competitive buyers when the prior belief is sufficiently diffuse and the news process sufficiently informative.

In addition to the intrinsic interest in the effect of buyer competition/market power, we believe this comparison sheds new light on the interpretation of the Coasian force. Often, the force is interpreted as: competition with the future self is sufficient to simulate the competitive outcome, and therefore leads to efficient trade. With news however, DG12 shows that the competitive outcome features periods of delay, and therefore is *not* efficient, even without adverse selection. We believe this suggests a slightly different interpretation of the Coasian force. Namely, the inability to commit to prices leads the buyer (i.e., uninformed party) to gain nothing from the ability to screen using prices. From this observation, it *follows* that trade will be efficient in the absence of adverse selection, regardless of news quality, and that it will simulate the competitive outcome without news as conjectured by Coase and substantiated by Fudenberg et al. (1985); Gul et al. (1986).

Our work is also related to Deneckere and Liang (2006) and Fuchs and Skrzypacz (2010) (hereafter, DL06 and FL10), both of which investigate frequent-offer, bilateral bargaining games. DL06 investigate an interdependent-value setting in the absence of news. They find that the presence of delay depends on the severity of adverse selection given the buyer's prior belief. During a period of delay, the buyer's belief must be exactly such that the

Coasian desire to speed up trade is absent, which is non-generic. The addition of learning via a diffusion process, even if arbitrarily noisy, means that the buyer's belief cannot remain constant over time at such a belief. As a result, our findings are considerably different from theirs even in the limit as the news becomes completely uninformative. FL10 study the independent-value setting from the Coase conjecture literature, with the addition of a Poisson arrival of a game-ending "event." The primary interpretation given to the event is the arrival of a new trader, but it can also be interpreted as arrival of a signal of the informed party's private information. The critical difference is that this information alters the support of the uninformed party's belief, unlike our Brownian news process. The possibility of this signal arrival delays trade, but their results are consistent with our interpretation of the Coasian force as evaporating the gains from screening via prices.

2 The Model

There are two players, a seller and a buyer, and a single durable asset of type $\theta \in \{L, H\}$, which is the seller's private information. Let $P_0 \in (0,1)$ denote the prior probability that the buyer assigns to $\theta = H$. The seller's (opportunity) cost of parting with the asset is K_θ, where we normalize $K_L = 0$, and the buyer's value for acquiring it is V_θ. Naturally, $K_H > 0$ and $V_H \geq V_L$. There is common knowledge of gains from trade: $V_\theta > K_\theta$ for each θ.

As discussed in the Introduction, bargaining dynamics will depend on whether or not a static adverse selection problem can arise. As in DG12, we define the condition as follows:

Definition 1. *The **Static Lemons Condition (SLC)** holds if and only if* $K_H > V_L$.[3]

Until Section 7, we assume the SLC holds.

The game is played in continuous time, starting at $t = 0$ with a potentially infinite horizon. At every time t, the buyer makes a price offer to the seller. If the seller accepts an offer of w at time t, the trade is executed and the game ends. The payoffs to the seller and the buyer respectively are $e^{-rt}(w - K_\theta)$ and $e^{-rt}(V_\theta - w)$, where $r > 0$ is the common discount rate. If trade never occurs, then both players receive a payoff of zero. Both players are risk-neutral and maximize their expected payoff.

In addition, until a trade agreement is reached, news about the seller's type is gradually revealed—as we discuss next. A heuristic description of the timing is depicted in Figure 1.

[3]The SLC is related to the *Static Incentive Constraint* of DL06, which is satisfied if and only if $K_H \leq E[V_\theta | P_0]$. Hence, the SLC implies that there exists at least *some* non-degenerate P_0 such that this Static Incentive Constraint fails.

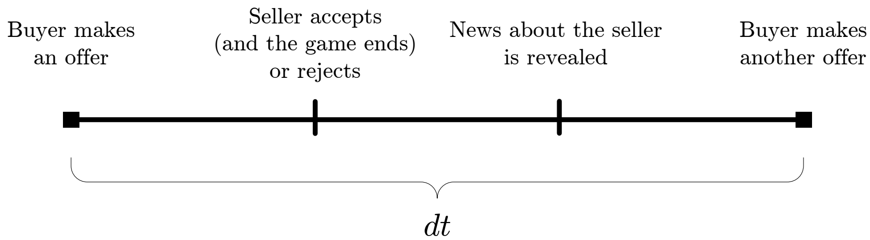

Figure 1: Heuristic Timeline of a Single "Period"

2.1 News Arrival

News about the seller's asset is revealed via a Brownian diffusion process. Regardless of type, the seller starts with an initial score X_0, normalized to 0. The news process then evolves according to

$$dX_t = \mu_\theta dt + \sigma dB_t \tag{1}$$

where $B = \{B_t, \mathcal{F}_t, 0 \le t \le \infty\}$ is standard Brownian motion on the canonical probability space $\{\Omega, \mathcal{F}, \mathcal{Q}\}$. At each time t, the entire history of news, $\{X_s, 0 \le s \le t\}$, is publicly observable. Without loss of generality, $\mu_H \ge \mu_L$. The parameters μ_H, μ_L and σ are common knowledge. Define the signal-to-noise ratio $\phi \equiv (\mu_H - \mu_L)/\sigma$. When $\phi = 0$, the news is completely uninformative. Larger values of ϕ imply more informative news. In what follows, we assume that $\phi > 0$, unless otherwise stated.

2.2 Equilibrium

Below we lay out the components of and requirements for equilibrium in turn, and collect them in Definition 2.

Stationarity

In keeping with the literature, we focus on behavior that is stationary, using the uninformed party's belief as the state variable. At every time t, if trade has not yet occurred, the buyer assigns a probability, $P_t \in [0,1]$, to $\theta = H$. Analytically, it is convenient to track the belief in terms of its log-likelihood ratio, denoted $Z_t \equiv \ln\left(\frac{P_t}{1-P_t}\right) \in \overline{\mathbb{R}}$ (i.e., the extended real numbers).[4] This transformation from belief as a probability to a log-likelihood ratio is injective, and therefore without loss. We will use z when referring to the state variable as opposed to the stochastic process Z (i.e., if $Z_t = z$, then the game is "in state z, at time t").

Formally, the belief process Z is adapted to the filtration $(\mathcal{H}_t)_{t\ge0}$, where $\mathcal{H}_t$ is the σ-algebra generated by $\{X_s, 0 \le s \le t\}$. X and Z are stochastic processes defined over the probability space $\{\Omega', \mathcal{H}, \mathcal{P}\}$, where $\Omega' = \Omega \times \Theta$, $\mathcal{H} = \mathcal{F} \times 2^\Theta$ and $\mathcal{P} = \mathcal{Q} \times \nu$, where ν is

[4]However, the degenerate beliefs $z = \pm\infty$ (i.e., $p = 0, 1$), are never reached in equilibrium and play a role in our analysis. Any reference to a generic state z should be intepreted as $z \in \mathbb{R}$ unless otherwise indicated.

the measure over $\Theta \equiv \{L, H\}$ defined implicitly by P_0.

Stationarity requires that both the current offer and the evolution of the belief depend only on the current belief.

Eq. Condition 1 (Stationarity)**.** *The buyer's offer in state z is given by $W(z)$, where $W : \overline{\mathbb{R}} \to \mathbb{R}$ is a Borel-measurable function, and Z is a time-homogenous $\mathcal{H}_t$-Markov process.*[5]

The Seller's Problem

The seller takes the offer function, W, as given. A pure strategy for the type-θ seller is then an $\mathcal{H}_t$-adapted stopping time $\tau_\theta(\omega) : \Omega' \to \mathbb{R}_+ \cup \{\infty\}$.[6] A mixed strategy for the seller is a distribution over such times, which can be represented as a stochastic process $S^\theta = \{S^\theta_t, 0 \leq t \leq \infty\}$ also adapted to $(\mathcal{H}_t)_{t \geq 0}$. The process must be right-continuous and satisfy $0 \leq S^\theta_t \leq S^\theta_{t'} \leq 1$ for all $t \leq t'$. $S^\theta(\omega)$ is a CDF over the type-θ seller's acceptance time on $\mathbb{R}_+ \cup \{\infty\}$ along the sample path $X(\omega, \theta)$. A discontinuous increase in S^θ corresponds to acceptance with an atom of probability mass.

Let $\mathcal{T}$ be the set of all $\mathcal{H}$-adapted stopping times. Given any offer function W and belief process Z, the type-θ seller faces a stopping problem.

$$\sup_{\tau \in \mathcal{T}} E^\theta \left[e^{-r\tau} (W(Z_\tau) - K_\theta) \right] \qquad (SP_\theta)$$

Recall that S^θ is a distribution over stopping times. Let $\mathcal{S}^\theta = supp(S^\theta)$. We say that S^θ solves (SP_θ) if all $\tau \in \mathcal{S}^\theta$ solve (SP_θ).[7]

Eq. Condition 2 (Seller Optimality)**.** *Given W and Z, S^θ solves the type-θ seller's problem (SP_θ).*

Consistent Beliefs

If trade has not occurred by time t, the buyer's belief, Z_t, is conditioned on both the entire path of past news and on the fact that the seller has rejected all past offers. It will be convenient to separate these two sources of information. Let f^θ_t denote the density of X_t

[5]This implies that Z is a time-homogenous Markov process with respect to the seller's information as well. For any t, s, because the distribution of Z_{t+s} given $\mathcal{H}_t$ depends only on Z_t, the distribution of Z_{t+s} given $\mathcal{H}_t$ and θ depends only on Z_t and θ, since $X(\cdot, \theta)$ has stationary, independent increments. In addition, while it is conventional to define stationarity as a restriction on strategies, which then has implications for beliefs through the *Belief Consistency* condition, Eq. Condition 1 is clearer in our model. That is, an alternative condition for *Stationarity* would replace the restriction on Z with a more notationally cumbersome, equivalent restriction on seller strategies.

[6]That is, τ_θ does not specify how to handle off-path offers, which is addressed by Eq. Condition 5.

[7]That is, for any $\tau_\theta \in \mathcal{S}^\theta$, $E^\theta \left[e^{-r\tau_\theta}(W(Z_{\tau_\theta}) - K_\theta)\right] = \sup_{\tau \in \mathcal{T}} E^\theta \left[e^{-r\tau}(W(Z_\tau) - K_\theta)\right]$.

conditional on θ, which for $t > 0$ is normally distributed with mean $\mu_\theta t$ and variance $\sigma^2 t$.[8] Let $S^\theta_{t^-} \equiv \lim_{s \uparrow t} S^\theta_s$ (which is well defined for $t > 0$ given that S^θ is bounded and non-decreasing), and specify that $S^\theta_{0^-} = 0$. Belief "at time t" should be interpreted to mean *before* observing the seller's decision at time t, which is why left limits are appropriate. If $S^L_{t^-} \cdot S^H_{t^-} < 1$ (i.e., given the sequence of offers up to t, there is positive probability that the seller will not have accepted yet in equilibrium), then the probability the buyer assigns to $\theta = H$ is defined by Bayes Rule as

$$\frac{P_0 f^H_t(X_t)(1 - S^H_{t^-})}{P_0 f^H_t(X_t)(1 - S^H_{t^-}) + (1 - P_0) f^L_t(X_t)(1 - S^L_{t^-})} \tag{2}$$

Taking the log-likelihood ratio of (2) results in

$$Z_t = \underbrace{\ln\left(\frac{P_0}{1 - P_0}\right) + \ln\left(\frac{f^H_t(X_t)}{f^L_t(X_t)}\right)}_{\hat{Z}_t} + \underbrace{\ln\left(\frac{1 - S^H_{t^-}}{1 - S^L_{t^-}}\right)}_{Q_t} \tag{3}$$

By working in log-likelihood space we are able to represent Bayesian updating as a linear process, and the buyer's belief as the sum of two components, $Z = \hat{Z} + Q$, as seen in (3). Notice that the two component processes separate the two sources of information to the buyer. $\hat{Z}$ is the belief process for a Bayesian who updates *only based on news* starting from $\hat{Z}_0 = Z_0 = \ln\left(\frac{P_0}{1-P_0}\right)$. Q is the stochastic process that keeps track of the information conveyed by the fact that the seller has rejected all past offers.

Eq. Condition 3 (Belief Consistency)**.** *For all t such that $S^L_{t^-} \cdot S^H_{t^-} < 1$, Z_t is given by (3).*

Option for Immediate Trade

The next condition is simple: if the buyer offers K_H, then both types accept with probability one. Since the buyer has all of the offering power, this feature is easy to establish in any discrete-time analog.[9]

Eq. Condition 4 (Option for Immediate Trade)**.** *If $W(Z_t) = K_H$, then $S^L_t = S^H_t = 1$.*

The Buyer's Problem

Given *Stationarity*, the value functions for each player depend only on the current state. Let $F_\theta(z)$ denote the expected payoff for the type-θ seller given state z. That is, for any $\tau \in \mathcal{S}^\theta$

[8] Let $f^H_0 = f^L_0$ be the Dirac delta function.

[9] See Fudenberg and Tirole (1991, pp. 409). Ortner (2016) imposes a similar condition in a continuous-time bargaining model.

$$F_\theta(z) \equiv E_z^\theta \left[e^{-r\tau}(W(Z_\tau) - K_\theta) \right],$$

where E_z^θ is the expectation with respect to the probability law of the process $\{Z_t, 0 \le t \le \infty\}$ conditional on θ and $Z_0 = z$. Similarly, let $F_B(z)$ denote the expected payoff to the seller in any given state z:

$$F_B(z) \equiv (1-p(z))E_z^L \left[\int_0^\infty (V_L - W(Z_t)) dS_{t^-}^L \right] + p(z) E_z^H \left[\int_0^\infty e^{-rt}(V_H - W(Z_t)) dS_{t^-}^H \right], \quad (4)$$

where $p(z) \equiv \frac{e^z}{1+e^z}$.

Taking the reservation values of each type seller as given, the buyer has essentially three options in any state z. He can make an offer of K_H and trade immediately. He can make a non-serious offer that both types will reject and wait for more news. Or, he can make an intermediate offer that will be rejected by the high type, but has positive probability of acceptance by the low type.

Rather than write the buyer's problem in terms of offers, it will be more convenient to do so in terms of "quantities" (i.e., the probability of trade).[10] Thus, the buyer's problem is to choose a stopping time, denoted by T, at which he trades for sure at price K_H and a process, denoted by Q, representing the intensity of trade with the low type prior to T. The intensity of trade at time $t < T$, dQ_t, determines the belief conditional on rejection (in accordance with (3) above), and therefore the price at time t must be the low type's expected payoff conditional on rejecting the offer (i.e., $W(Z_t) = F_L(Z_t) = F_L(\hat{Z}_t + Q_t)$). We refer to the pair (T, Q) as a *policy*. A policy is *feasible* if T is an $\mathcal{F}_t$-measurable stopping rule and Q is non-negative, non-decreasing process, $\mathcal{F}_t$-measurable process. Let Γ denote the set of feasible policies.

Eq. Condition 5 (Buyer Optimality)**.** *For any z, F_B as defined by* (4) *satisfies:*

$$F_B(z) = \sup_{(Q,T)\in\Gamma} \left\{ (1-p(z))E_z^L \left[\int_0^T e^{-rt}(V_L - F_L(\hat{Z}_t + Q_t))e^{-Q_{t^-}} dQ_t + e^{-rT}e^{-Q_T}(V_L - K_H) \right] + p(z)E_z^H \left[e^{-rT}(V_H - K_H) \right] \right\} \quad (5)$$

Definition 2. *An **equilibrium** of the model is a quadruple (W, S^L, S^H, Z) that satisfies Equilibrium Conditions 1-5.*

[10] Formally dealing with continuation play following deviations from W posses well-known existence problems in a continuous-time setting (Simon and Stinchcombe, 1989). And would require a substantially more complicated set of available strategies for the seller.

3 Equilibrium Construction

In this section we construct our equilibrium of interest. It is characterized by a belief threshold, β, and, for all $z < \beta$, a rate of trade with the low type. Specifically, when $z \geq \beta$, the buyer offers $W(z) = K_H$, which is accepted by both types of the seller, meaning trade is immediate. When $z < \beta$, the buyer offers some $W(z) < K_H$, which the high type rejects for sure. The low type accepts at a state-specific flow rate (i.e., proportional to time), meaning rejection is a (weakly) positive signal that $\theta = H$. Therefore, the consistent belief conditional on rejection, Z, has additional upward drift, denoted $\dot{q}(z) \geq 0$.

If $\dot{q}(z) > 0$, then the low type is accepting at a strictly positive flow rate, and must be indifferent between accepting and rejecting. The offer, then, must equal what the low type would earn by rejecting. Alternatively, $\dot{q}(z) = 0$ would correspond to no trade and the belief would update only based on news. For this case, it is without loss to specify the offer be the payoff the low type would earn by rejecting. However, we will demonstrate that $\dot{q}(z) = 0$ cannot happen in equilibrium.

The next definition gives a formal description of the equilibrium candidate parameterized by $(\beta, \dot{q})$.

Definition 3. *For $\beta \in \overline{\mathbb{R}}$ and measurable, absolutely continuous function $\dot{q} : (-\infty, \beta) \to \mathbb{R}_+$, let $T(\beta) \equiv \inf\{t : Z_t \geq \beta\}$ and $\Sigma(\beta, \dot{q})$ be the strategy profile and belief process:*

$$Z_t = \begin{cases} \hat{Z}_t + \int_0^t \dot{q}(Z_s)ds & \text{if } t \leq T(\beta) \\ \text{arbitrary} & \text{otherwise}^{11} \end{cases} \tag{6}$$

$$S_t^H = \begin{cases} 0 & \text{if } t < T(\beta) \\ 1 & \text{otherwise} \end{cases} \tag{7}$$

$$S_t^L = \begin{cases} 1 - e^{-\int_0^t \dot{q}(Z_s)ds} & \text{if } t < T(\beta) \\ 1 & \text{otherwise} \end{cases} \tag{8}$$

$$W(z) = \begin{cases} K_H & \text{if } z \geq \beta \\ E_z^L[e^{-rT(\beta)}]K_H & \text{if } z < \beta \end{cases} \tag{9}$$

Theorem 1. *There exists a unique pair $(\beta, \dot{q})$ such that $\Sigma(\beta, \dot{q})$ is an equilibrium. The unique β and $\dot{q}$ are given in Lemmas 1 and 3 below.*

[11] According to $\Sigma(\beta, \dot{q})$, if $t > T(\beta)$, trade should commence by time t with probability one. Hence, the evolution of Z—the belief conditional on rejection—in this event is the specification of the buyer's off-path beliefs. Because the buyer never offers more than K_H, no matter how high Z becomes, the specification of these off-path belief has no bearing on our results (in contrast to DG12 and others, where off-path beliefs matter, and a mild refinement on them is imposed).

The theorem is established by construction. Below we demonstrate how necessary conditions identify the unique $(\beta, \dot{q})$-pair. Proofs of the lemmas and the theorem are found in the Appendix.

3.1 Necessary Conditions: Determining β (and F_B)

Let $\bar{V}(z) \equiv E_z[V_\theta]$. For any state $z \geq \beta$, the buyer's value is $F_B(z) = \bar{V}(z) - K_H$. For any $z < \beta$ the buyer's value satisfies

$$rF_B(z) = \frac{\dot{q}(z)}{1+e^z}(V_L - F_L(z) - F_B(z)) + \left(\frac{\phi^2}{2}(2p(z)-1) + \dot{q}(z)\right) F_B'(z) + \frac{\phi^2}{2} F_B''(z). \quad (10)$$

Collecting the $\dot{q}$ terms gives,

$$rF_B(z) = \frac{\phi^2}{2}(2p(z)-1) F_B'(z) + \frac{\phi^2}{2} F_B''(z) + \dot{q}(z)\left(\frac{1}{1+e^z}\big(V_L - F_L(z) - F_B(z)\big) + F_B'(z)\right). \quad (11)$$

Let $U(z, z') \equiv \frac{p(z')-p(z)}{p(z')}(V_L - F_L(z')) + \frac{p(z)}{p(z')}F_B(z')$, which represents the buyer's payoff from moving the (post-rejection) belief to z' starting from state $z \leq z'$. In a Σ-equilibrium the belief does not jump, meaning $z' = z$ must be weakly optimal. The associated first-order condition is

$$U_2(z, z) = \frac{1}{1+e^z}(V_L - F_L(z) - F_B(z)) + F_B'(z) \leq 0. \quad (12)$$

So either, $U_2(z, z) = 0$ or $U_2(z, z) < 0$. But if $U_2(z, z) < 0$ then, the buyer strictly prefers *not* to move the belief up, implying $\dot{q}(z) = 0$. Hence, in either case,

$$\dot{q}(z)\left(\frac{1}{1+e^z}\big(V_L - F_L(z) - F_B(z)\big) + F_B'\right) = 0. \quad (13)$$

Therefore, (11) simplifies to

$$rF_B(z) = \frac{\phi^2}{2}(2p(z)-1) F_B'(z) + \frac{\phi^2}{2} F_B''(z). \quad (14)$$

The ODE has unique closed-form solution

$$F_B(z) = \frac{1}{1+e^z} C_1 e^{u_1 z} + \frac{1}{1+e^z} C_2 e^{u_2 z}, \quad (15)$$

where $(u_1, u_2) = \frac{1}{2}(1 \pm \sqrt{1 + 8r/\phi^2})$ and C_1, C_2 are two constants to be determined by two boundary conditions.[12] The boundary conditions are:

$$\lim_{z\to-\infty} F_B(z) < \infty \tag{16}$$

$$F_B(\beta) = \bar{V}(\beta) - K_H. \tag{17}$$

Because the buyer's payoff is uniformly bounded by V_H, (16) must hold (which implies $C_2 = 0$). When Z_t hits β, trade is immediate regardless of θ. Hence, (17) is the required *value-matching* condition.

Finally, *smooth pasting* of F_B is required at β:

$$F_B'(\beta) = \bar{V}'(\beta). \tag{18}$$

To see why smooth pasting is required, consider the buyer at $z = \beta$. Given (17), if $F_B'(\beta^-) < \bar{V}'(\beta)$, then a convex combination of $F_B(\beta - \epsilon)$ and $\bar{V}(\beta + \epsilon) - K_H$ is strictly greater then $F_B(\beta) = \bar{V}(\beta) - K_H$. This implies that the buyer can improve his payoff by simply waiting (i.e., make non-serious offers) for all $z \in [\beta, \beta + \delta)$ for sufficiently small δ. On the other hand, if $F_B'(\beta^-) > \bar{V}'(\beta)$, then there exists an ϵ such that $F_B(\beta - \epsilon) < \bar{V}(\beta - \epsilon) - K_H$, meaning the buyer would do better to trade at K_H immediately, in violation of Eq. Conditions 4-5.[13]

These necessary conditions yield a unique solution for β and F_B. Let $\underline{z} \equiv \ln\left(\frac{K_H - V_L}{V_H - K_H}\right)$. That is, $\underline{z}$ is the unique solution to $\bar{V}(z) = K_H$.

Lemma 1. *If $\Sigma(\beta, \dot{q})$ is an equilibrium, then $\beta = \beta^* \equiv \underline{z} + \ln\left(\frac{u_1}{u_1 - 1}\right)$. In addition, for all $z \geq \beta$, $F_B(z) = \bar{V}(z) - K_H$, and for all $z < \beta$, $F_B(z)$ is given by (15), with $C_1 = C_1^* \equiv \frac{K_H - V_L}{u_1 - 1}\left(\frac{u_1}{u_1 - 1}\frac{K_H - V_L}{V_H - K_H}\right)^{-u_1}$ and $C_2 = C_2^* \equiv 0$.*

3.2 Necessary Conditions: Determining $\dot{q}$ (and F_L)

First, in the candidate, the low type weakly prefers to reject $W(z)$ when $z < \beta$. Hence, her equilibrium payoff must be equal to the payoff she would obtain by always rejecting in these states, and waiting for K_H to be offered: for all z, $F_L(z) = E_z^L[e^{-rT(\beta)}]K_H$. So, for $z \geq \beta$, $F_L(z) = K_H$. For $z < \beta$, F_L must satisfy

$$rF_L(z) = \left(\dot{q}(z) - \frac{\phi^2}{2}\right) F_L'(z) + \frac{\phi^2}{2} F_L''(z). \tag{19}$$

[12] Polyanin and Zaitsev (2003, p. 215)

[13] See Shiryaev (1978, Sect. 3.8) for a more formal treatment of the necessity of smooth-pasting conditions or Dixit (1993, Sect. 4.1) for a more intuitive exposition.

Solving for $\dot{q}(z)$ gives that

$$\dot{q}(z) = \frac{rF_L(z) + \frac{\phi^2}{2}F_L'(z) - \frac{\phi^2}{2}F_L''(z)}{F_L'(z)}. \tag{20}$$

Second, recall that (12) gave $U_2(z,z) = \frac{1}{1+e^z}\left(V_L - F_L(z) - F_B(z)\right) + F_B'(z) \leq 0$, meaning the buyer weakly prefers not to move the belief up slightly through "faster" trade with the low type. The next lemma states that the buyer is indifferent toward this option, meaning (12) holds with equality.

Briefly, if $U_2(z,z) < 0$, then by (13), $\dot{q}(z) = 0$. This lack of additional drift means it takes longer for Z_t to reach β, negatively impacting the low type's continuation value, which (we just argued) coincides with F_L. This, in turn, raises $U_2(z,z)$ and leads to a violation of (12). The interpretation is that, if trade ever came to a halt, the low type's continuation value would become so low that she would be too cheap for the buyer not to trade with.

Lemma 2. *If $\Sigma(\beta,\dot{q})$ is an equilibrium, then for all $z < \beta$,*

$$U_2(z,z) = \frac{1}{1+e^z}\big(V_L - F_L(z) - F_B(z)\big) + F_B'(z) = 0.$$

Rearranging, then using Lemma 1's characterization of F_B, gives that, for all $z < \beta$,

$$F_L(z) = (1+e^z)F_B'(z) + V_L - F_B(z) \tag{21}$$

$$= V_L + C_1^*(u_1 - 1)e^{u_1 z} \tag{22}$$

Substituting (22) into (20) gives

$$\dot{q}(z) = \frac{rV_L e^{-u_1 z}}{C_1^* u_1(u_1 - 1)} > 0 \tag{23}$$

Lemma 3. *If $\Sigma(\beta,\dot{q})$ is an equilibrium, then $\beta = \beta^*$ and, for all $z < \beta^*$, $\dot{q}(z)$ is given by* (23)*, and $F_L(z)$ is given by* (22)*. Immediately, $F_L(z) = K_H$ for all $z \geq \beta^*$.*

Henceforth, we use $(\beta,\dot{q})$ in reference to the pair that characterize the unique Σ-equilibrium as stated in Theorem 1.

4 Discussion of Bargaining Dynamics

The intensity of trade is characterized by β and $\dot{q}$. If $z \geq \beta$, the buyer puts sufficient probability that $\theta = H$, such that he simply ends the game by offering K_H inducing immediate trade. If $z < \beta$, he does not yet believe $\theta = H$ is sufficiently likely, so instead makes offers that only the low type may find acceptable. The induced intensity of trade with the low type is captured by $\dot{q}(z)$, so is proportional to time, and the associated price is precisely the low type's continuation value, $W(z) = F_L(z)$. Specifically, this rate of trade is bounded away from zero, strictly decreasing, and grows infinitely large as the buyer becomes more and more convinced that $\theta = L$.

Property 1. *In the unique Σ-equilibrium, $\dot{q}$ is continuous, $\dot{q}(z) \gg 0$, and $\dot{q}'(z) < 0$ for all $z < \beta$. Further, $\dot{q}(z) \to \infty$ as $z \to -\infty$ (i.e., as $p \to 0$).*[14]

Conditional on no trade, then, the buyer's belief drifts up compared to the inference he would make based on the realization of the news alone. This force is strongest when z is low and decreases as z increases toward β. Because price equals low-type continuation value, it continuously increases with z below β.

Property 2. *In the unique Σ-equilibrium, W is continuous and $W'(z) > 0$ for all $z < \beta$. Further, $W(z) \to V_L$ as $z \to -\infty$ (i.e., as $p \to 0$).*

Perhaps surprisingly, the price offer given any (non-degenerate) belief is strictly greater than V_L, even for beliefs where only the low type trades (i.e., $z < \beta$). That is, if $z < \beta$, then there is positive probability of trade before the belief reaches β, but conditional on trade, θ is revealed to be L meaning the buyer earns a negative payoff. Conditional on trade below β, then, the buyer always experiences "ex-post regret."[15]

Making these relatively high offers can be rationalized as a form of costly experimentation. The buyer's value function is increasing, meaning he values pushing the belief up. Making an offer that the low type may accept, generates a potential benefit (rejection raises the belief and, therefore, the buyer's expected payoff), but also a potential cost (acceptance means the buyer overpays, and earns a negative payoff). In equilibrium, these forces exactly balance out due to the buyer's inability to commit to future offers, as we discuss next.

4.1 The Buyer's Payoff and the Coasian Force

In equilibrium, the buyer makes offers that only the low type might accept when $z < \beta$. Consider a simplified version of the model in which the buyer cannot do so. Specifically, at every time t, the buyer is restricted to either make an offer of K_H, which will be accepted by both types with probability one, or make no offer.

This "no-price-screening" model (NPS, hereafter) essentially reduces to a standard optimal stopping problem for the buyer. His belief updates only based on news, $Z = \hat{Z}$, and

[14] Note, all three statements also hold for the unconditional rate of trade: $(1 - p(z))\dot{q}(z)$.

[15] Unlike in DL06 and FS10.

stopping corresponds to offering K_H. If he stops when his belief is $Z_t = z$, his payoff is $\bar{V}(z) - K_H$. Hence, he chooses an optimal stopping time, T, to maximize $E[e^{-rT}(V_\theta - K_H)]$.

It is not difficult to establish that the NPS model has a unique solution, which is a threshold policy: $T_{NPS} = \inf\{t : Z_t \geq \beta_{NPS}\}$, where $-\infty \leq \beta_{NPS} < \infty$. Further, below β_{NPS}, the buyer's value function satisfies the ODE in (14) in the NPS model. Finally, the value-matching and smooth-pasting conditions (16)-(18) are also required. Therefore,

Property 3. *$\beta = \beta_{NPS}$, and the buyer's payoff in the unique Σ-equilibrium, F_B, is identical to his payoff in the unique solution to the NPS model.*

Returning to our game, intuition might have suggested that the buyer will make use of the news in two ways: *i*) to ensure he is sufficiently confident that $\theta = H$, before offering the price needed to compensate a H-type seller, and *ii*) to extract value out of the L-type seller with low prices if he becomes sufficiently confident that $\theta = L$. Our result is consistent with (*i*), but not (*ii*). Even though the buyer does offer prices offer than K_H if his belief is below β^*, and there is probability that such a price will be accepted, the buyer's equilibrium payoff, F_B is still identical to what he would garner if he had no ability to screen using prices. This can be viewed as the manifestation of the "Coasian" force in our model.

Starting for a low belief, the buyer would like to be able commit to a low offer for at least some discrete interval of time. The rejection of this offer would, however, increase the buyer's belief at which point he would again be tempted to "experiment" by offering a price that may be accepted by the low type as described above. Without any ability to commit, he will indeed make this offer, which the low type foresees. This raising low-type continuation value, which coincides with price.

Corollary 1. *In the unique Σ-equilibrium, F_B is continuous, $F_B'(z) > 0$ for all z, and $F_B(z) \to 0$ as $z \to -\infty$ (i.e., as $p \to 0$).*

As the belief that $\theta = L$ becomes arbitrarily strong, the buyer's value in the game vanishes. However, this is *not* due to destruction of total surplus through inefficient delay. As noted above, the rate of rate with the low type grows arbitrarily large, and the low-type seller's value tends to V_L as $z \to -\infty$. Hence, trade is efficient in this limit (discussed further below), but all of the surplus is captured by the low type.

4.2 Efficiency

In the absence trade, each player gets a payoff of zero. The (expected) *surplus* obtained by the seller's side of the game in state z is given by

$$\Pi^S(z) \equiv (1 - p(z))(F_L(z) - K_L) + p(z)(F_H(z) - K_H).$$

The buyer's suplus in state z is simply $F_B(z)$. So, total surplus realized in state z is then given by $\Pi(z) \equiv \Pi^S(z) + F_B(z)$. Due to common knowledge of gains from trade, the efficient outcome is to trade immediately, resulting in a total potential (or first-best) surplus of

$$\Pi^{FB}(z) \equiv (1 - p(z))(V_L - K_L) + p(z)(V_H - K_H).$$

Hence, $\Pi^{FB}(z) - \Pi(z) \geq 0$ and any strictly positive difference is the efficiency loss due to expected delays in trade. We define the normalized loss in efficiency as a function of z by

$$\mathcal{L}(z) \equiv \frac{\Pi^{FB}(z) - \Pi(z)}{\Pi^{FB}(z)}.$$

Corollary 2. *In the unique* Σ*-equilibrium,* $\mathcal{L}$ *is continous,* $\mathcal{L}(z) = 0$ *if and only if* $z \geq \beta$*, but* $\mathcal{L}(z) \to 0$ *as* $z \to -\infty$ *(i.e., as* $p \to 0$*).*

5 The Effect of Buyer Competition

In this section, we explore the effect of competition among buyers by contrasting our findings with DG12, which analyzes an analogous setting except with perfectly competitive buyers.[16] In most economic settings, one expects a more competitive market to lead to more efficient outcomes. However, when the uninformed side of the market can learn from news, we will see that introducing competition can have exactly the opposite effect.

By way of terminology, we refer to the *competitive outcome* as the equilibrium with multiple competing buyers from DG12, and the *bilateral outcome* as the unique Σ-equilibrium with only a single buyer. Notionally, we use a subscript $s \in \{b(ilateral), c(competetive)\}$ on objects when referencing the respective outcomes.

When buyers are competitive (and the SLC holds), DG12 show that the unique equilibrium is characterized by a pair of beliefs $\alpha_c < \beta_c$ and the following three regions. For $z > \beta_c$, trade takes place immediately at a price $\bar{V}(z)$. For $z < \alpha_c$, buyers offer V_L, the high type rejects and the low type mixes. Conditional on a rejection at some $z < \alpha_c$, buyer's belief jumps to α_c. For all $z \in (\alpha_c, \beta_c)$ trade occurs with probability zero and the buyers' beliefs evolve solely due to news.

[16] Specifically, they replace the *Option for Immediate Trade* and *Buyer Optimality* equilibrium conditions with *Zero Profit* and *No Deals* conditions. The first ensures that any trade that takes place earns zero expected profit for a buyer. The second ensures that when trade does not take place, there does not exist an offer that a buyer could make and earn positive profit. They also impose a modest refinement on off-path beliefs.

Bilateral outcome	Competitive outcome
Trade is efficient for $z \geq \beta_b$. For $z < \beta_b$: - Probability of trade is strictly positive, and rate is proportional to time. - Rate is decreasing in z.	Trade is efficient for $z \geq \beta_c$. For $z < \beta_c$: - Zero probability of trade for $z \in (\alpha_c, \beta_c)$. - Atom of trade for $z < \alpha_c$.

TABLE 1: Comparison of bilateral and competetive outcomes when the SLC holds.

Table 1 compares the competitive outcome to the bilateral one. Both outcomes involve a threshold belief above which trade is fully efficient and below which there is positive probability of delay. However, unlike the smooth trading probability in the bilateral outcome, the trading probability below the threshold in the competitive outcome is lumpy (either zero or an atom). Given the upper threshold determines the set of states in which the outcome is fully efficient, the following lemma has important efficiency implications.

Lemma 4. $\beta_c > \beta_b$.

The intuition behind this lemma is the following. In the competitive outcome, buyers are willing to offer $\bar{V}(z)$ at any z such that the high-type seller is willing to accept. Thus, it is the high-type seller that decides when to "stop," which nets her $\bar{V}(z) - K_H$. In the bilateral outcome, it is the buyer who decides when to "stop" (i.e., offer K_H) which nets him $\bar{V}(z) - K_H$ (recall that β_b is the same as the solution to the buyer's stopping problem when he is unable to screen). While the net payoff to player who determines when to stop in the respective settings is the same, they have different expectations about the evolution of $\hat{Z}$. In particular, the drift of $\hat{Z}$ under the high-type seller's filtration is strictly greater than under the buyer's filtration. Hence, the solution to the high-type's stopping problem involves waiting longer (i.e., a higher threshold). The intuition above is further strengthened by the lower boundary, α_c, in the competitive outcome whereat the low-type seller mixing "pushes" the belief process upward, making the high-type even more willing to wait.

Clearly, Lemma 4 implies there exists a set of states (i.e., $z \in (\beta_b, \beta_c)$) such that the bilateral outcome is fully efficient and the competitive outcome is not. By continuity, the bilateral outcome remains more efficient just below β_b. However, for low z, the ranking reverses and the competitive outcome is more efficient as can be seen in Figure 2 and in the following proposition.[17]

Proposition 1. *There exist a $z_1 < z_2$, both in $(-\infty, \beta_b)$, such that,*

- $\mathcal{L}_c(z) \geq \mathcal{L}_b(z)$ *for all $z > z_2$ where the inequality is strict for all $z \in (z_2, \beta_c)$, and*

[17] Thorough the figures, beliefs are measured as probabilities with $b_b = p(\beta_b)$, $b_c = p(\beta_c)$, and so forth.

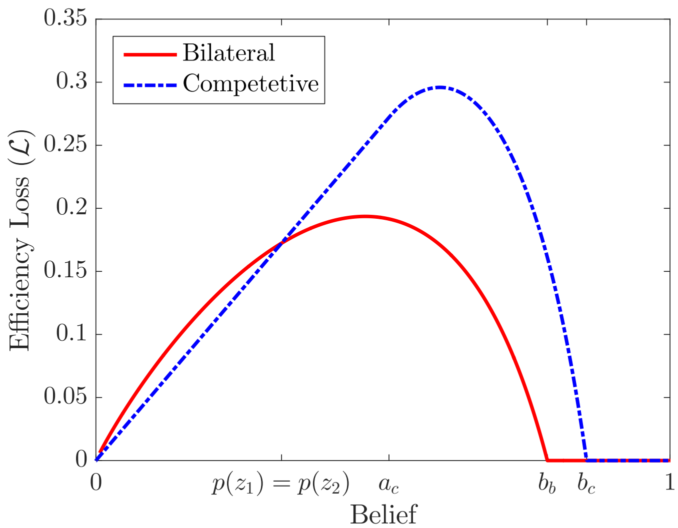

FIGURE 2: Comparison of loss in efficiency across the bilateral (solid) and competitive (dashed) outcomes.

- $\mathcal{L}_c(z) < \mathcal{L}_b(z)$ *for all* $z < z_1$.

Intuitively, when the belief is low, trade is more efficient in the competitive outcome because the low type is trading more rapidly (i.e., with an atom compared to with a rate in the bilateral outcome), and when z is low it is the low type's trading behavior that determines efficiency.[18]

In terms of player welfare, the comparison for the both the buyer and the high-type seller is trivial. The buyer earns positive surplus (for all z) in the bilateral outcome, and zero in the competitive one. Conversely, the high-type seller earns zero surplus in the bilateral outcome (since the price never exceeds K_H), but earns positive surplus in the competitive outcome.

The comparison for the low-type seller is more nuanced. When the belief is low, she is better off in the bilateral outcome than in the competitive, but the reverse when the belief is high, as seen in Figure 3. As discussed in Section 4, in the bilateral setting, the buyer offers prices above V_L as form of experimentation. This benefits the low-type seller. There is no scope for this costly experimentation in the competitive setting, as buyer-profits are driven to zero. In contrast, when the belief is high, the low-type seller enjoys buyer competition since it raises the price to $\bar{V}$ instead of only K_H.

[18]In Figure 2, $z_1 = z_2$. This feature appears to be general, but is yet unproven.

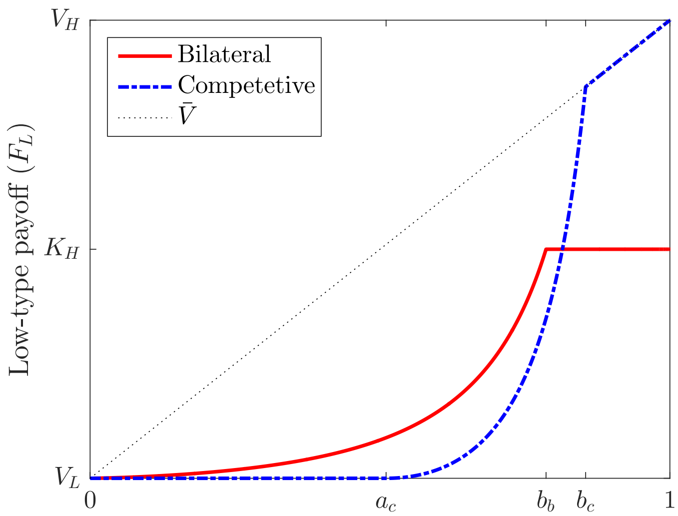

FIGURE 3: Comparison of the low-type seller payoff across the bilateral (solid red line) and competitive (dashed blue line) outcomes.

6 Effect of News Quality

In this section we investigate the effect of news quality. First, we explore how an increase in news quality affects equilibrium play and payoffs. Then we take the limit as news becomes arbitrarily informative (i.e., $\phi \to \infty$) and arbitrarily noisy (i.e., $\phi \to 0$). Finally, we compare the $\phi \to 0$ limit to a model with no news analyzed by DL06.

6.1 An increase in news quality

We first state the formal results and then discuss their intuition.

Proposition 2. *In the unique Σ-equilibrium, as the quality of news, ϕ, increases:*

(i) β *increases.*

(ii) The rate of trade, $\dot{q}$, *decreases for* $z < \beta - \frac{2u_1-1}{u_1(u_1-1)}$ *but increases for* $z \in (\beta - \frac{2u_1-1}{u_1(u_1-1)}, \beta)$.

(iii) The buyer's payoff increases for all $z < \beta$.

(iv) The low-type seller's payoff increases for $z < \beta - \frac{1}{u_1-1}$ *but decreases for* $z \in (\beta - \frac{1}{u_1-1}, \beta)$.

(v) Total surplus increases for $z < \beta - \frac{1}{u_1}$ *but decreases for* $z \in (\beta - \frac{1}{u_1}, \beta)$.

Intuitively, as the quality of news increases, the buyer learns about the seller's type faster, and therefore finds it optimal to choose a higher belief threshold before exercising the option for immediate trade. Thus, both β and F_B increase with ϕ. These findings are

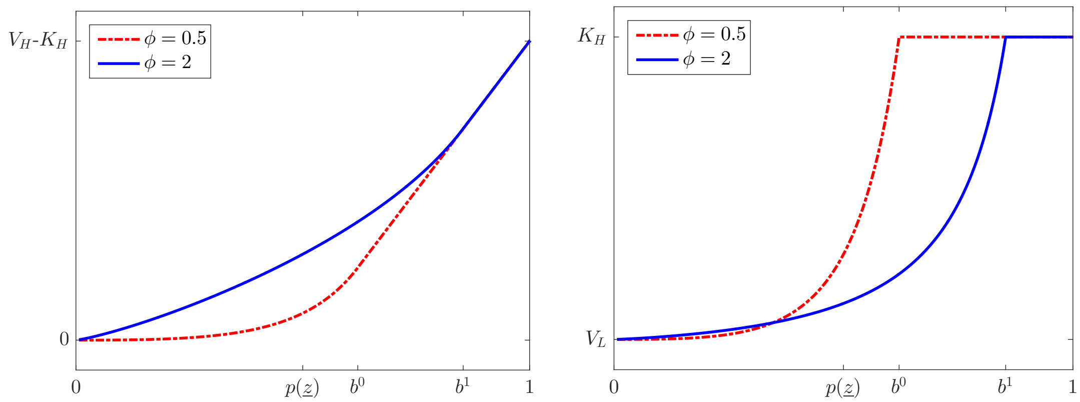

(a) Buyer Payoff (F_B) (b) Low Type Payoff (F_L)

FIGURE 4: The effect of news quality on equilibrium payoffs.

illustrated in Figure 4(a). The effect of news quality on F_L is more subtle because there are several opposing forces. To understand them, recall that the low type's equilibrium payoff is equal to the expected discounted value of waiting until $z = \beta$, when K_H is offered. Now, holding β and $\dot{q}$ fixed, a higher ϕ means an increase in the volatility of $\hat{Z}$ which reduces the expected waiting cost and therefore increases F_L. On the other hand, a higher β (or lower $\dot{q}$) increases the waiting costs, thereby decreasing F_L. For intuition about (iii), consider a discrete increase in news quality from ϕ^0 to ϕ^1 and therefore by (i), $\beta^0 < \beta^1$. Clearly, the low type must be worse off with the higher news quality for $z \in (\beta^0, \beta^1)$. Continuity implies this ranking must persist for z just below β^0. However, for low enough z, the volatility effect dominates as illustrated in Figure 4(b).

These same opposing forces also affect the overall efficiency as illustrated in Figure 5(a). On the one hand, a higher ϕ "speeds things up" and reduces $\mathcal{L}$. On the other hand, because β increases, there are states in which trade would be fully efficient under ϕ^0, but is delayed with positive probability under ϕ^1. Thus, a higher ϕ leads to less efficient outcomes for z near the upper threshold, while the first effect dominates and $\mathcal{L}$ decreases for low z.

6.2 Arbitrarily informative news ($\phi \to \infty$)

The following proposition characterizes the limit properties of the equilibrium as news quality becomes arbitrarily high. Let $\overset{pw}{\to}$ and $\overset{u}{\to}$ denote pointwise and uniform convergence, respectively.

Proposition 3. *In the unique Σ-equilibrium, as $\phi \to \infty$:*

(i) $\beta \to \infty$.

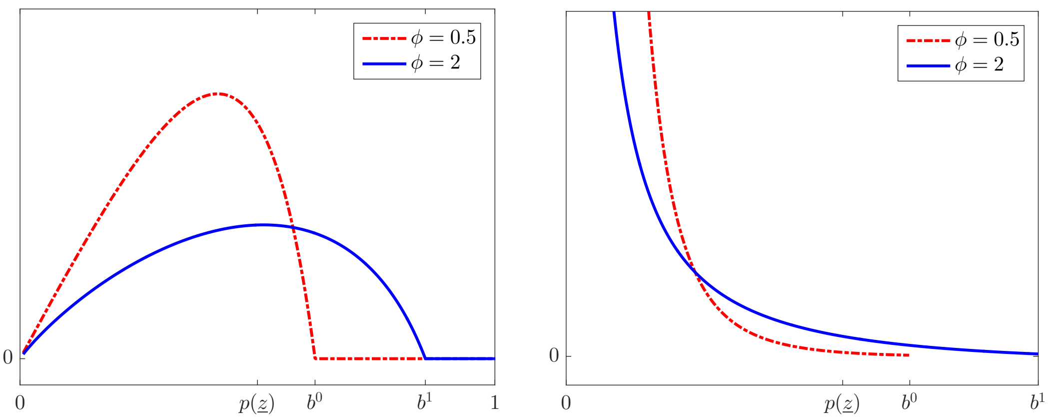

(a) Efficiency loss ($\mathcal{L}$)

(b) Rate of trade ($\dot{q}$)

FIGURE 5: The effect of news quality on inefficiency and the rate of trade.

(ii) $\dot{q} \stackrel{pw}{\to} \infty$, *but for any* $x > 0$, $\dot{q}(\beta - x) \to \frac{rV_L}{K_H - V_L} e^x$.

(iii) $F_B \stackrel{u}{\to} p(z)(V_H - K_H)$.

(iv) $F_L \stackrel{pw}{\to} V_L$.

(v) $\mathcal{L} \stackrel{u}{\to} 0$.

Property (i) says that the buyer waits until she is virtually sure that the seller is a high type before offering K_H. As $\phi \to \infty$, this learning happens so quickly that the delay becomes trivial and the buyer captures all of the surplus from trading with the high type seller (i.e., $V_H - K_H$). Intuition might suggest that a similar type of pattern should obtain when trading with a low type. That is, one might expect the buyer would wait until she is virtually sure that the seller is of the low type before offering K_L; this learning would happen arbitrarily quickly as $\phi \to \infty$; and thereby the buyer would also extract all the surplus from trading the low-type seller.

Recall from Section 4.1, however, that this intuition is incorrect. For *any* ϕ, as $z \to -\infty$: $\dot{q}(z) \to \infty$, $F_B(z) \to 0$, $F_L(z) \to V_L$ and $\mathcal{L}(z) \to 0$, due to the Coasian force. Properties (ii)-(v) demonstrate that this temptation to speed up trade with the low type overwhelms the motivation to learn about the seller's type, even when this learning takes place arbitrarily quickly.[19]

Properties (iii)-(v) are illustrated in Figure 6. The disparity between the strength of convergence for F_L and F_B is due to the fact that, even for large ϕ, $F_L(z) = K_H$ for all

[19]This fact may be partially due to the order of limits. By analyzing a continuous-time model directly, we have implicitly taken the period length to zero first (i.e., before taking $\phi \to \infty$). If we were to interchange the order of limits (i.e., consider a discrete-time model with news and take the limit as $\phi \to \infty$ *before* taking the period length to zero), then it is plausible that the intuition given above would prove correct.

$z \geq \beta$, meaning the convergence of F_L to V_L is only pointwise.

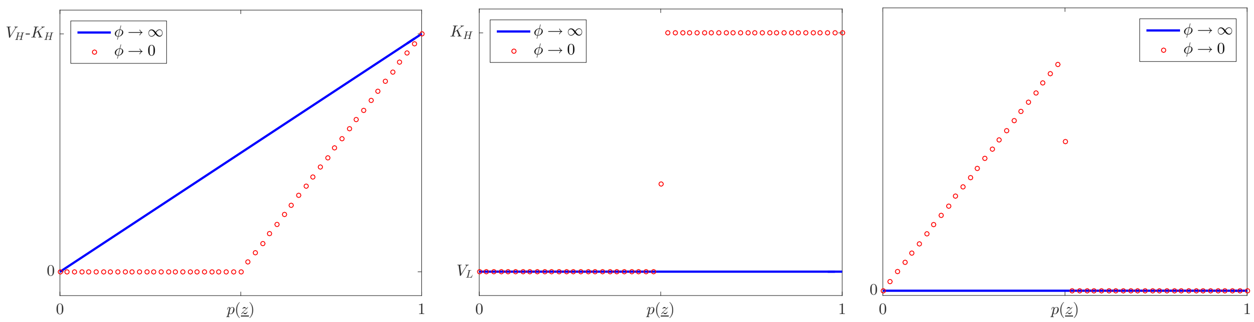

(a) Buyer Payoff (F_B) (b) Low Type Payoff (F_L) (c) Efficiency loss ($\mathcal{L}$)

FIGURE 6: Limiting payoffs and efficiency loss as $\phi \to \infty$ and $\phi \to 0$.

6.3 Arbitrarily uninformative news ($\phi \to 0$)

We now turn to the other extreme in which news tends to pure noise.

Proposition 4. *In the unique Σ-equilibrium, as $\phi \to 0$:*

(i) $\beta \to \underline{z}$.

(ii) For all $z < \underline{z}$, $\dot{q}(z) \to \infty$, *but* $\dot{q}(\underline{z}) \to 0$.

(iii) $$F_B \xrightarrow{u} \begin{cases} 0 & \text{if } z < \underline{z} \\ \bar{V}(z) - K_H & \text{if } z \geq \underline{z}. \end{cases}$$

(iv) $$F_L \xrightarrow{pw} \begin{cases} V_L & \text{if } z < \underline{z} \\ (1-e^{-1})V_L + e^{-1}K_H & \text{if } z = \underline{z} \\ K_H & \text{if } z > \underline{z}. \end{cases}$$

(v) $$\mathcal{L} \xrightarrow{pw} \begin{cases} \frac{p(z)(V_H - K_H)}{\Pi^{FB}(z)} & \text{if } z < \underline{z} \\ \frac{p(z)(V_H-K_H)-(1-p(z))e^{-1}(K_H-V_L)}{\Pi^{FB}(z)} & \text{if } z = \underline{z} \\ 0 & \text{if } z > \underline{z}. \end{cases}$$

To interpret these results, it is useful to draw a comparison to DL06. For convenience, we restate their result below using our notation.

Result (DL06, Proposition 2)**.** *Consider a two-type, discrete-time model with no news (i.e., $\phi = 0$), and suppose that SLC holds. In equilibrium, as the period length between offers goes to zero,*

(a) *For all* $z > \underline{z}$*, the buyer offers* K_H *and the seller accepts w.p.1.*

(b) *For* $z < \underline{z}$*, the buyer makes an offer of* $w_0 = \frac{V_L^2}{C_H}$*. The high type rejects and the low type mixes such that the belief is* $\underline{z}$ *following a rejection.*

(c) *For* $z = \underline{z}$*, there is delay of length* 2τ*, where* τ *satisfies* $V_L = e^{-r\tau} K_H$*, after which the buyer offers* K_H *and the seller accepts w.p.1.*

There are notable similarities between the result above and our findings in Proposition 4. For $z > \underline{z}$, the predictions are perfectly aligned; trade takes place immediately at a price equal to the high-type's cost. In addition, for $z < \underline{z}$, in both settings there is a "burst" of trade with the low type and delay ensues conditional on a rejection. The key differences are the buyer's offer when $z < \underline{z}$ and the amount of ensuing delay. In our case, the offer is V_L and the amount of ensuing delay is τ, whereas in DL06 the offer is $w_0 < V_L$ and the amount of ensuing delay is exactly twice as along.

A (perhaps) surprising implication is that the buyer is strictly worse off for all $z < \underline{z}$ in our limit (continuous time, $\phi \to 0$) than in that of DL06 (discrete time, $\phi = 0$, period length $\to 0$). The intuition has two components. First, without news, if the buyer delays trade (by making unacceptable offers), the belief remains constant. Second, when the buyer's belief is $\underline{z}$, the temptation to speed up trade (i.e., the Coasian force) is absent because the buyer's continuation value from this state is zero (in the no-news, continuous-time limit). Hence, without news, the buyer can leverage an endogenous form of commitment power: it is both feasible and sequentially rational for the buyer to delay trade keeping the belief at $\underline{z}$. This allows him to extract more surplus from the low type in states $z < \underline{z}$.

In contrast, with even an arbitrarily small amount of Brownian news, the buyer's belief will instantaneously diverge from $\underline{z}$ almost surely. That is, the buyer cannot "sit" at $\underline{z}$, and make non-serious offers for any amount of time, because he observes news and updates his belief arbitrarily quickly, which strengthens the Coasian force and reduces his ability to extract surplus in all states $z < \underline{z}$.

Another implication is that even a small amount of news can lead to a discontinuous improvement in efficiency. Without news, in order to extract the extra surplus, the buyer uses his (endogenous) commitment power at $\underline{z}$, which implies more delay and hence more inefficiency. These findings are illustrated in Figure 7.

7 When the SLC Fails and the Coase Conjecture

We now turn to equilibrium when the SLC fails. In this case, there is no trade delay in the unique Σ-equilibrium.

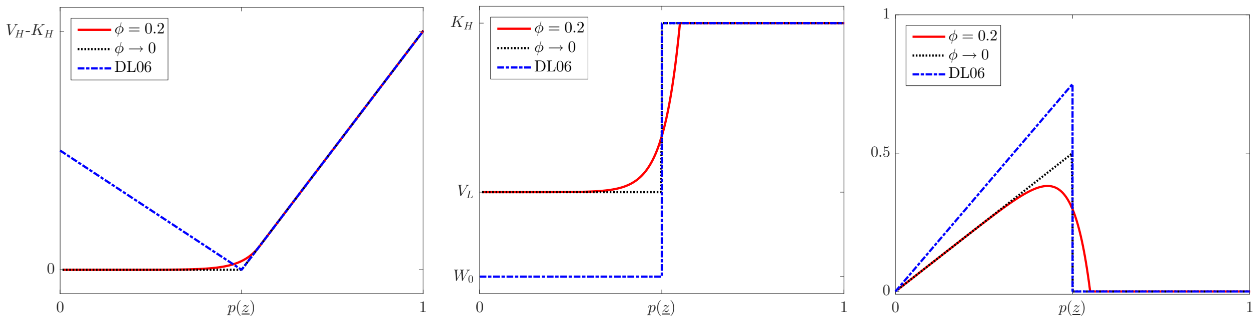

(a) Buyer Payoff (F_B) (b) Low Type Payoff (F_L) (c) Efficiency loss ($\mathcal{L}$)

FIGURE 7: Payoffs and efficiency comparison to DL06 with $\phi > 0$ and as $\phi \to 0$.

Theorem 2. *When the SLC fails, in the unique Σ-equilibrium, $\beta = -\infty$ (i.e., $W(z) = K_H$ and trade is immediate for all z).*

Verification of the equilibrium is straightforward. The uniqueness result is obtained by following the same steps from Section 3, and finding that no $\beta \in \mathbb{R}$ can satisfy the necessary condition (15)-(18).

One intuition for the result comes via the connection to the NPS model from Section 4. Again the buyer's equilibrium payoff (in the true game) coincides with his payoff in the solution to the NPS model. Without the SLC, however, the solution to the NPS model is to "stop" (i.e., trade at price K_H) immediately. Why? Measuring beliefs as probabilities, $p \in (0, 1)$, in the NPS model, the buyer's reward from stopping is positive and linear in his belief, his belief is a martingale, and he discounts payoffs. Hence, he can do not better than stopping immediately.

Strikingly, Theorem 2 holds regardless of the quality of the news process, ϕ. This can be viewed as an extension of Coase's conjecture. Within our environment, Coase conjectured that the buyer's competition with his future self would lead to immediate trade at price K_H when there is no news, $\phi = 0$, and independent values, $V_H = V_L > K_H$ (which is a special case of the SLC failing). Our results show that this force swamps any incentive to delay and wait from Brownian news.

However, the result brings a subtlety to the interpretation of the Coasian force. Often, the force is interpreted as: competition with the future self is sufficient to simulate the competitive outcome, and therefore leads to efficient trade. With news however, DG12 shows that competitive outcome features periods of delay, and therefore is *not* efficient, even when the SLC fails (Proposition 5.3 therein).

We believe this suggests a slightly different interpretation of the Coasian force. Namely, the inability to commit to prices leads the buyer (i.e., uninformed party) to gain nothing from the ability to screen using prices. From this observation, it *follows* that *i*) trade will be efficient when the SLC fails, regardless of ϕ, and *ii*) it will simulate the competitive outcome when $\phi = 0$. However, neither of these features need hold in general (e.g., they both fail to hold in our model under the SLC).

8 Concluding Remarks

We have investigated a bilateral-bargaining model in which the seller's private information is gradually revealed to the buyer until agreement is reached. In our equilibrium of interest, the buyer benefits from the presence of this news process only to the degree that he would benefit from it in an environment that prohibited screening with prices (as captured by the NPS). In addition, both the equilibrium bargaining dynamics and efficiency differ from the competitive-buyer analog. Hence, insofar as the buyer "competes with his future self," as is commonly said in similar frequent-offer bargaining settings, this inter-temporal competition is *not* a perfect proxy for intra-temporal competition.

We also demonstrated that this equilibrium is unique among candidates that feature a threshold for efficient trade and trade proportional to time below the threshold, a seemingly natural form. We conjecture, however, that this equilibrium is unique even when not restricting to this form. One way to prove the conjecture involves two key steps, which are currently in progress. First, that any non-Σ equilibrium would have to consist of both atoms of trade for some beliefs and intervals of beliefs with zero probability of trade (reminiscent of DG12). Second, that such a no-trade interval is not possible, by generalizing the analysis of Section 3.2 and Lemma 2. The interpretation again being that if trade ever came to a halt, the low type's continuation value would become so low that she would be too cheap for the buyer not to trade with.

References

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

Daley, B. and B. Green (2012). Waiting for news in the market for lemons. *Econometrica 80*, 1433–1504.

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

Dixit, A. (1993). *The Art of Smooth Pasting*. Harwood Academic Publishers.

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

Fudenberg, D., D. K. Levine, and J. Tirole (1985). Infinite-horizon models of bargaining with one-sided incomplete information. *Game-Theoretic Models of Bargaining*, 73–98.

Fudenberg, D. and J. Tirole (1991). *Game Theory*. MIT Press.

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

Harrison, M. J. (2013). *Brownian Models of Performance and Control*. Cambridge University Press.

Ortner, J. (2016). Durable goods monopoly with stochastic costs. *Theoretical Economics, Forthcoming*.

Polyanin, A. D. and V. F. Zaitsev (2003). *Handbook of Exact Solutions for Ordinary Differential Equations*. Chapman & Hall.

Shiryaev, A. N. (1978). *Optimal Stopping Rules*. Springer.

Simon, L. K. and M. B. Stinchcombe (1989). Extensive form games in continuous time: Pure strategies. *Econometrica 57*, 1171–1214.

Wang, X. (2015). Catering innovation: Entrepreneurship and the acquisition market. *Working Paper*.

A Appendix

Proof of Lemma 1. From Section 3.1, if $\beta \in \mathbb{R}$, then F_B, β, C_1, C_2 must satisfy (15)-(18). First, from (15),

$$\lim_{z\to-\infty} F_B(z) = \lim_{z\to-\infty} \frac{1}{1+e^z} C_1 e^{u_1 z} + \frac{1}{1+e^z} C_2 e^{u_2 z} = \begin{cases} -\infty & \textit{if } C_2 < 0 \\ 0 & \textit{if } C_2 = 0 \\ \infty & \textit{if } C_2 > 0. \end{cases}$$

To satisfy (16), therefore, in any solution $C_2 = 0$. This simplifies the remaining two equations, (17) and (18):

$$F_B(\beta) = \frac{C_1 e^{u_1\beta}}{1+e^\beta} = \bar{V}(\beta) - K_H = \frac{e^\beta}{1+e^\beta}(V_H - V_L) + V_L - K_H$$
$$F_B'(\beta) = \frac{C_1 e^{u_1\beta}\left((u_1-1)e^\beta + u_1\right)}{(1+e^\beta)^2} = \bar{V}'(\beta) = \frac{e^\beta}{(1+e^\beta)^2}(V_H - V_L).$$

The unique solution to the pair is given by

$$\beta = \beta^* \equiv \ln\left(\frac{K_H - V_L}{V_H - K_H}\right) + \ln\left(\frac{u_1}{u_1 - 1}\right)$$
$$C_1 = C_1^* \equiv \frac{K_H - V_L}{u_1 - 1}\left(\frac{u_1}{u_1-1}\frac{K_H - V_L}{V_H - K_H}\right)^{-u_1}.$$

If $\beta = \infty$, then $F_B(z) = 0$ for all $z \in \mathbb{R}$. But then the buyer would improve his payoff by offering K_H (leading to payoff $\bar{V}(z) - K_H$) for any $z > \underline{z}$. Finally, if $\beta = -\infty$, then $F_B(z) = \bar{V}(z) - K_H$ for all $z \in \mathbb{R}$. But then the buyer's payoff is negative all $z < \underline{z}$, and he would improve his payoff by making a non-serious offer in these states. □

Proof of Lemma 2. Fix $\beta = \beta^*$ and F_B as given by Lemma 1. Given an arbitrary (absolutely continuous) $\dot{q}$ on $z < \beta$, let $G_L^{\dot{q}}(z)$ be the expected payoff of a low type who rejects all offers until $Z_t \geq \beta$ (i.e., $E_z^L[e^{-rT(\beta)}]$). Let $\dot{q}^*$ denote expression for $\dot{q}$ given in (23). Therefore, for all $z < \beta$,

$$\frac{1}{1+e^z}\left(V_L - G_L^{\dot{q}^*}(z) - F_B(z)\right) + F_B'(z) = 0.$$

From Section 3.1, $U_2(z,z) \leq 0$ for all $z < \beta$. For the purpose of contradiction, suppose there exists $z_0 < z_1 < \beta$ such that $U_2(z,z) < 0$ for all z in (z_0, z_1). To satisfy (13), then $\dot{q}(z) = 0$ for all $z \in (z_0, z_1)$. Hence, $G_L^{\dot{q}}(z) < G_L^{\dot{q}^*}(z)$ for all $z \in (z_0, z_1)$. This implies that, for all $z \in (z_0, z_1)$,

$$\frac{1}{1+e^z}\left(V_L - G_L^{\dot{q}}(z) - F_B(z)\right) + F_B'(z) > 0.$$

Finally, recall that in equilibrium, the low type always weakly prefers rejection in state $z < \beta$, so $F_L(z) = G_L^{\dot{q}}(z)$. Hence, for all $z \in (z_0, z_1)$,

$$U_2(z,z) = \frac{1}{1+e^z}\left(V_L - F_L(z) - F_B(z)\right) + F_B'(z) > 0,$$

producing a contradiction. Finally, since for any (absolutely continuous) $\dot{q}$, $G_L^{\dot{q}}$ and F_B are continuous on $(-\infty, \beta)$, and in equilibrium $F_L(z) = G_L^{\dot{q}}(z)$, $U_2(z,z)$ is continuous. Hence, if $\Sigma(\beta, \dot{q})$ is an equilibrium, $U_2(z,z) = 0$ for all $z < \beta$. □

Proof of Lemma 3. Immediate from Lemmas 1 and 2, and analysis in Section 3.2. □

Proof of Theorem 1. Lemmas 1 and 3 show that there exists a unique candidate $\Sigma(\beta, \dot{q})$. Thus, to prove the theorem, we need only verify that this candidate satisfies the equilibrium conditions. Eq. Conditions 1, 3, and 4 are satisfied by construction for any $(\beta, \dot{q})$: 1 follows immediately from (6), 3 can be verified by inserting (7) and (8) into (3), 4 is also immediate from (7)-(9) since $S_t^\theta = 1$ for all $t \geq T(\beta) = \inf\{s : Z_s \geq \beta\} = \inf\{s : W(Z_s) = K_H\}$.

Next we verify Eq. Condition 2 (seller optimality). Consider first the high type and note from (7) that $\mathcal{S}^H = \{T(\beta)\}$ and from (9) that $W(z) \leq K_H$. Therefore,

$$\sup_{\tau \in \mathcal{T}} E^H \left[e^{-r\tau}(W(Z_\tau) - K_H) \right] \leq 0 = F_H(z),$$

where $F_H(z)$ is equal to the high-type's payoff under the candidate equilibrium strategy, $T(\beta)$, which verifies that S^H solves (SP_H).

For the low type, recall that, by construction, $F_L(z) = E_z^L[e^{-rT(\beta)}]K_H$. Let $\mathcal{T}(\beta) \equiv \mathcal{T} \cap \{\tau : \tau \leq T(\beta), \forall\, \omega\}$, i.e., the set of all $\mathcal{H}$-adapted stopping times such that $\tau \leq T(\beta)$ for all ω. Observe that $E_z^L[e^{-r\tau}W(Z_\tau)] \leq F_L(z)$ for any $\tau \in \mathcal{T} \setminus T(\beta)$ since W is bounded above by K_H and delay is costly. That is, since K_H is the largest possible offer, it is optimal for the low type to accept it as soon as it is offered. Note further that $\mathcal{S}^L \subseteq \mathcal{T}(\beta)$. To prove S^L solves (SP_L), we show that, in fact, for *any* $\tau \in \mathcal{T}(\beta)$, $E_z^L\left[e^{-r\tau}W(Z_\tau)\right] = F_L(z)$, which verifies that S^L solves (SP_L).

Let $f_L(t,z) \equiv e^{-rt}W(z)$ and note that f_L is C^2 for all $z \neq \beta$. Conditional on $\theta = L$ and $t < T(\beta)$, Z evolves according to

$$dZ_t = \left(\dot{q}(Z_t) - \frac{\phi^2}{2}\right) dt + \phi dB_t.$$

By Dynkin's formula, for any $\tau \in \mathcal{T}(\beta)$

$$E_z^L[f_L(\tau, Z_\tau)] = f_L(0,z) + E_z^L\left[\int_0^\tau \mathcal{A}^L f_L(s, Z_s) ds\right],$$

where $\mathcal{A}^L$ is the characteristic operator for the process $Y_t = (t, Z_t)$ under $\mathcal{Q}^L$, i.e.,

$$\mathcal{A}^L f(t,z) = \frac{\partial f}{\partial t} + \left(\dot{q}(z) - \frac{\phi^2}{2}\right)\frac{\partial f}{\partial z} + \frac{1}{2}\phi^2 \frac{\partial^2 f}{\partial z^2}. \tag{24}$$

Applying $\mathcal{A}^L$ to f_L, we get that

$$\begin{aligned}\mathcal{A}^L f_L(t,z) &= e^{-rt}\left[-rW(z) + \left(\dot{q}(z) - \frac{\phi^2}{2}\right)W'(z) + \frac{\phi^2}{2}W''(z)\right] \\ &= e^{-rt}\left[-rF_L(z) + \left(\dot{q}(z) - \frac{\phi^2}{2}\right)F_L'(z) + \frac{\phi^2}{2}F_L''(z)\right] \\ &= 0,\end{aligned}$$

where the first equality follows from the fact the $W(z) = F_L(z)$ (by construction, see (9)) and the second equality from the fact that $\dot{q}$ satisfies (20). Hence, for any $\tau \in \mathcal{T}(\beta)$, $E_z^L[f_L(\tau, Z_\tau)] = F_L(z)$, as desired.

The last step in the proof is to verify Eq. Condition 5 (buyer optimality). In order to do so, we first characterize an upper bound on the buyer's payoff in Lemma A.1 (below) and then verify that F_B achieves this bound. An immediate corollary of Lemma A.1 is that if there exists a feasible (Q,T) under which the buyer's expected payoff satisfies the hypothesis of the Lemma, then the policy is optimal. By construction, F_B is the buyer's payoff under the policy $Q_t = \int_0^t \dot{q}(Z_s)ds$, $T = T(\beta)$. Observe that $F_B \in C^1$ and is C^2 for all $z \neq \beta$, therefore, it suffices to verify that F_B satisfies (A.1)-(A.3).

Verification that F_B satisfies (A.1)-(A.3)*:*

- For $z \leq \beta$. First, note that (A.2) holds with equality for all $z < \beta$ by construction. Hence, we need only verify (A.1) and (A.3). For (A.1), define

$$K(z,z') \equiv \frac{p(z') - p(z)}{p(z')}(V_L - F_L(z')) + \frac{p(z)}{p(z')}F_B(z')$$

  To see that $F_B(z) \geq \sup_{z' \geq z}\{K(z,z')\}$, note that

$$\frac{d}{dz'}K(z,z') = \frac{C_1 e^{(u_1-1)z'}\left(e^z - e^{z'}\right)(-1+u_1)u_1}{1+e^z} < 0,\ \forall z' \in (z,\beta).$$

  Since $K(z,z')$ is decreasing in z', we have that $F_B(z) = K(z,z) = \sup_{z' \in (z,\beta)} K(z,z')$. Furthermore, $K(z,z') = \bar{V}(z) - K_H \leq F_B(z)$ for $z' \geq \beta$ (as shown below), which verifies that the first term is non-positive for all $z' > z$.

  To see that $F_B \geq \bar{V} - K_H$, note that

$$F_B(\beta - x) - (\bar{V}(\beta - x) - K_H) = \frac{e^{-u_1 x}\left(e^x + e^{x(1+u_1)}(u_1 - 1) - u_1 e^{u_1 x}\right)(V_H - K_H)(K_H - V_L)}{e^x(u_1 - 1)(V_H - K_H) + u_1(K_H - V_L)}$$

  The denominator on the RHS is positive since $V_H > K_H > V_L$. The numerator is positive provided that for all $x > 0$, $e^x + e^{x(1+u_1)}(u_1 - 1) - u_1 e^{u_1 x} \geq 0$, which can be shown to hold for all $u_1 \geq 1$ (i.e., over the entire relevant parameter space).

- For $z > \beta$. First, note that $F_B = \bar{V} - K_H$ by construction so (A.3) holds with equality. Hence, it remains to verify (A.1) and (A.2). Since $F_L(z) = K_H$ for all $z \geq \beta$, we get that

$$
\begin{aligned}
&\frac{p(z')-p(z)}{p(z')}(V_L - F_L(z')) + \frac{p(z)}{p(z')}F_B(z') \\
&\qquad\qquad\qquad\qquad = \frac{p(z')-p(z)}{p(z')}\big(V_L - K_H\big) + \frac{p(z)}{p(z')}\big(\bar{V}(z') - K_H\big) \\
&\qquad\qquad\qquad\qquad = \frac{p(z')-p(z)}{p(z')}\big(V_L - K_H\big) + \frac{p(z)}{p(z')}\big(p(z')V_H + (1-p(z'))V_L - K_H\big) \\
&\qquad\qquad\qquad\qquad = p(z)V_H + (1-p(z)V_L - K_H \\
&\qquad\qquad\qquad\qquad = \bar{V}(z) - K_H,
\end{aligned}
$$

and therefore (A.1) holds with equality *for all* $z' \geq z$. Verifying (A.2) is equivalent to showing that for all $z > \beta$,

$$
\frac{\phi^2}{2}\Big((2p(z)-1)\bar{V}'(z) + \bar{V}''(z)\Big) - r(\bar{V}(z) - K_H) \leq 0.
$$

Noting that $(2p(z)-1)\bar{V}'(z) + \bar{V}''(z) = 0$ and $\beta > \underline{z} \Rightarrow \bar{V}(z) - K_H > 0$ for all $z > \beta$ implies the above inequality and completes the proof.

□

Lemma A.1. *Let $F^{Q,T}(z)$ denote the buyer's payoff under an arbitrary feasible policy $(Q,T) \in \Gamma$ starting from $Z_0 = z$. Let $\mathcal{A}$ denote the characteristic operator of $\hat{Z}_t$ under $\mathcal{Q}^B$. Suppose that $f \in C^1$, $f \in C^2$ almost everywhere and satisfies*

$$
f(z) \geq \frac{p(z')-p(z)}{p(z')}(V_L - F_L(z')) + \frac{p(z)}{p(z')}f(z') \quad \text{for all } z' \geq z \in \mathbb{R}; \tag{A.1}
$$

$$
rf(z) \geq \mathcal{A}f(z) \quad \text{for almost all } z \in \mathbb{R}; \tag{A.2}
$$

$$
f(z) \geq V(z) - K_H \quad \text{for all } z \in \mathbb{R}, \tag{A.3}
$$

then $f \geq F^{Q,T}$.

Proof Sketch: If f is the buyer's value function, (A.1) says that the buyer cannot benefit by enforcing a jump from z to z' and is a standard optimality condition in impulse control. The inequality in (A.2) says that the buyer cannot benefit by making a non-serious offer and "wait for news" and is a standard optimality condition in optimal stopping. Condition (A.3) says that the buyer cannot benefit by stopping immediately (i.e., offering K_H). That (A.1)-(A.3) combined with the smoothness properties are sufficient for an upper bound on the buyer's payoff follows closely standard arguments (see e.g., Harrison (2013), Corollary 5.2, Proposition 7.2) and is therefore omitted. □

Proof of Lemma 4. As shown in DG12 (see the proof of Lemma B.3 therein), $\beta_c > z_H^*$, where z_H^* is the threshold belief at which a high-type seller would stop in a game where $\bar{V}(z)$ is always offered and beliefs evolve only according to news. Using the closed form expressions for z_H^* (see (41) in DG12) and β_b (see Lemma 1), it is straightforward to check that $z_H^* > \beta_b$, which proves the lemma. □

Proof of Proposition 1. First, $\mathcal{L}_b, \mathcal{L}_c \geq 0$, $\mathcal{L}_b(z) > 0$ if and only $z > \beta_b$, and $\mathcal{L}_c(z) > 0$ if and only $z > \beta_c$. By Lemma 4, $\beta_b < \beta_c$. Hence, by continuity of $\mathcal{L}_c$ and $\mathcal{L}_b$, there exists $z_2 < \beta_b$ such that $\mathcal{L}_b(z) > \mathcal{L}_c(z)$ for all $z \in (z_1, \beta_b)$.

In the bilateral outcome, $F_H^b = 0$, so $\Pi_b(z) = F_B^b(z) + (1 - p(z))\, F_L^b(z)$. In the competitive outcome, $F_B^c = 0$, so $\Pi_c(z) = p(z)F_H^c(z) + (1 - p(z))F_L^c(z)$. Further, in competitive outcome, for all $z < \alpha_c$, both seller payoffs are constant: $F_L^c(z) = V_L$ and $F_H^c(z) = A \in (K_H, V_H)$. Direct calculations then show:

$$\lim_{z\to-\infty} \mathcal{L}_b(z) = \lim_{z\to-\infty} \mathcal{L}_c(z) = 0.$$

Therefore, by L'Hospital's rule:

$$\lim_{z\to-\infty} \left(\frac{\mathcal{L}_b(z)}{\mathcal{L}_c(z)}\right) = \lim_{z\to-\infty} \left(\frac{\mathcal{L}_b'(z)}{\mathcal{L}_c'(z)}\right) = \frac{V_H - K_H}{V_H - K_H - A} > 1.$$

Hence, there exists $z_1 > -\infty$ such that $\mathcal{L}_b(z) > \mathcal{L}_c(z)$ for all $z < z_1$. □

Proof of Proposition 2. From the expression in Lemma 1, β is decreasing in u_1, which recall is defined as $u_1 \equiv \frac{1}{2}\left(1 + \sqrt{1 + 8r/\phi}\right)$. Clearly u_1 decreases with ϕ, which implies (i). For (ii), using the expression in (23) we have that

$$\frac{d}{du_1}\dot{q}(z) = \frac{rV_L}{e^{u_1 z}(u_1 - 1)^2 u_1^2 (K_H - V_L)} \zeta^{u_1} \left(1 + u_1(z - 2) - u_1^2 z + (u_1 - 1)u_1 \ln(\zeta)\right)$$

The expression above is strictly positive (negative) for $z > (<)\beta - \frac{2u_1 - 1}{u_1(u_1 - 1)}$, which implies (ii). For (iii), it is sufficient to show that F_B is decreasing in u_1 below β. To do so, plug in the expression for $C_1 = C_1^*$ into F_B and differentiate with respect to u_1 to get that

$$\begin{aligned}
\frac{d}{du_1}F_B(z) &= \frac{1}{1 + e^z} e^{u_1 z} \left(\frac{\partial C_1^*}{\partial u_1} + zC_1^*\right) \\
&= \frac{1}{1 + e^z} e^{u_1 z} \left(\frac{K_H - V_L}{u_1 - 1}\right) \zeta^{-u_1}(z - \ln(\zeta)) \\
&< 0,
\end{aligned}$$

where $\zeta \equiv \frac{u_1(K_H - V_L)}{(u_1 - 1)(V_H - K_H)} > 0$ and the inequality follows from noting that $\ln(\zeta) = \beta$. For (iv), note that for $z < \beta$,

$$\begin{aligned}
\frac{d}{du_1}F_L(z) &= e^{u_1 z} \left((1 + (u_1 - 1)z)C_1^* + (u_1 - 1)\frac{\partial C_1^*}{\partial u_1}\right) \\
&= e^{u_1 z} \left(\frac{K_H - V_L}{u_1 - 1}\right) \zeta^{-u_1} \left(1 + (u_1 - 1)(z - \ln(\zeta)\right).
\end{aligned}$$

Noting that $e^{u_1 z}\left(\frac{K_H - V_L}{u_1 - 1}\right)\zeta^{-u_1} > 0$, we have that $F_L(z)$ increases with u_1 (decreases with ϕ) for $z \in (\beta - \frac{1}{u_1 - 1}, \beta)$ and decreases in u_1 (increasing in ϕ) for $z < \beta - \frac{1}{u_1 - 1}$, which proves (iv). For (v), note that $\Pi(z) = F_B(z) + (1 - p(z))F_L(z)$ and therefore

$$\begin{aligned}\frac{d}{du_1}\Pi(z) &= \frac{d}{du_1}F_B(z) + (1 - p(z))\frac{d}{du_1}F_L(z)\\ &= \frac{1}{1+e^z}e^{u_1 z}\left(\frac{K_H - V_L}{u_1 - 1}\right)\zeta^{-u_1}\left(1 + u_1(z - \ln(\zeta))\right)\end{aligned}$$

Thus, Π increases with u_1 (decreases with ϕ) for $z \in (\beta - \frac{1}{u_1}, \beta)$ and decreases with u_1 (increases with ϕ) for $z < \beta - \frac{1}{u_1}$. As a result, (v) immediately follows.

□

Proof of Proposition 3. First, note that taking the limit as $\phi \to \infty$ is equivalent to taking the limit as $u_1 \to 1$ from above. For (i), using the expression for β in Lemma 1, we have that

$$\lim_{u_1 \to 1} \beta = \underline{z} + \lim_{u_1 \to 1} \ln\left(\frac{u_1}{u_1 - 1}\right) = \infty.$$

For (ii), using the expressions for $\dot{q}$ and C_1^* from Lemmas 3 and 1,

$$\dot{q}(z) = \frac{rV_L e^{-u_1 z}}{C_1^* u_1(u_1 - 1)} = \frac{rV_L e^{-u_1 z}\left(\frac{u_1(K_H - V_L)}{(u_1-1)(V_H - K_H)}\right)^{u_1}}{u_1(K_H - V_L)},$$

which, for all $z < \beta$, tends to ∞ as $u_1 \to 1$ from above. Incorporating the expression for β yields:

$$\dot{q}(\beta - x) = \frac{rV_L e^{u_1 x}}{u_1(K_H - V_L)} \to \frac{rV_L e^x}{K_H - V_L}$$

as u_1 goes to 1. For (iii), from Lemma 1,

$$F_B(z) = \begin{cases} \bar{V}(z) - K_H & \textit{if } z \geq \beta \\ \frac{e^{u_1 z}(V_H - K_H)\left(\frac{u_1(K_H - V_L)}{(u_1-1)(V_H - K_H)}\right)^{1-u_1}}{(1+e^z)u_1} & \textit{if } z < \beta \end{cases}$$

As $u_1 \to 1$, $\beta \to \infty$, meaning for any $z \in \mathbb{R}$,

$$\lim_{u_1 \to 1} F_B(z) = \lim_{u_1 \to 1} \frac{e^{u_1 z}(V_H - K_H)\left(\frac{u_1(K_H - V_L)}{(u_1-1)(V_H - K_H)}\right)^{1-u_1}}{(1+e^z)u_1} = \frac{e^z}{1+e^z}(V_H - K_H) = p(z)(V_H - K_H).$$

Further, since $F_B(z)$ is continuous in z and non-decreasing in ϕ (Proposition 2), the convergence is uniform by Dini's Theorem.[20] For (iv), from Lemma 3,

$$F_L(z) = \begin{cases} K_H & \textit{if } \ z \geq \beta \\ V_L + e^{u_1 z}(V_H - K_H)^{u_1}(K_H - V_L)\left(\frac{u_1(K_H - V_L)}{u_1 - 1}\right)^{-u_1} & \textit{if } \ z < \beta \end{cases}$$

As $u_1 \to 1$, $\beta \to \infty$, meaning for any $z \in \mathbb{R}$,

$$\lim_{u_1 \to 1} F_L(z) = V_L + \lim_{u_1 \to 1} e^{u_1 z}(V_H - K_H)^{u_1}(K_H - V_L)\left(\frac{u_1(K_H - V_L)}{u_1 - 1}\right)^{-u_1} = V_L.$$

Finally, for (v),

$$\begin{aligned} 0 \leq \mathcal{L}(z) = \frac{\Pi^{FB}(z) - \Pi(z)}{\Pi^{FB}(z)} &= \frac{p(z)(V_H - K_H) - F_B(z) + (1 - p(z))(V_L - F_L(z))}{\Pi^{FB}(z)} \\ &\leq \frac{p(z)(V_H - K_H) - F_B(z)}{\Pi^{FB}(z)}, \end{aligned} \tag{25}$$

where the last inequality follows from $F_L(z) \geq V_L$ for all z (regardless of ϕ). By (iii), the term in (25) uniformly converges to 0 as $u_1 \to 1$, implying $\mathcal{L}$ does as well. □

Proof of Proposition 4. First, note that taking the limit as $\phi \to 0$ is equivalent to taking the limit as $u_1 \to \infty$. For (i), using the expression for β in Lemma 1, we have that

$$\lim_{u_1 \to \infty} \beta = \underline{z} + \ln\left(\lim_{u_1 \to \infty} \frac{u_1}{u_1 - 1}\right) = \underline{z} + \ln(1) = \underline{z}.$$

From (23), we have that $\dot{q}(z) = \frac{rV_L}{C_1^* u_1(u_1-1)e^{u_1 z}}$. Therefore, to prove (ii) it suffices to show that $\lim_{u_1 \to \infty} C_1^* u_1(u_1 - 1)e^{u_1 z} = 0$ for $z < \underline{z}$ and $\lim_{u_1 \to \infty} C_1^* u_1(u_1 - 1)e^{u_1 \underline{z}} = \infty$. Using the closed form expression for C_1^* in Lemma 1, we have that

$$C_1^* u_1(u_1 - 1)e^{u_1 z} = (K_H - V_L)\left(\frac{u_1 - 1}{u_1}\right)^{u_1}\left(\frac{V_H - K_H}{K_H - V_L}e^z\right)^{u_1} u_1$$

The first term on the right hand side is a constant. The second term limits to e^{-1} as $u_1 \to \infty$. Thus, the remaining terms determine the limiting properties. They can be written as $u_1 y^{u_1}$, where $y \equiv \frac{V_H - K_H}{K_H - V_L}e^z$. Notice that $z < \underline{z} \Rightarrow y < 1 \Rightarrow u_1 y^{u_1} \to 0$, whereas $z = \underline{z} \Rightarrow y = 1 \Rightarrow u_1 y^{u_1} = u_1 \to \infty$. This completes the proof of (ii).

For (iii), note that for all $z \leq \underline{z}$, $0 \leq F_B(z) \leq C_1^* e^{u_1 z} \leq C_1^* e^{u_1 \underline{z}}$. And further, $C_1^* e^{u_1 \underline{z}} = (K_H - V_L)\left(\frac{u_1 - 1}{u_1}\right)^{u_1} \frac{1}{u_1 - 1} \to 0$ as $u_1 \to \infty$. Thus, we have obtained uniform bound on $F_B(z)$ below $\underline{z}$, which converges to zero implying the first part of (iii). That $F_B(z) \xrightarrow{u} \bar{V}(z) - K_H$ for $z \geq \underline{z}$ follows from continuity of F_B, $F_B(z) = \bar{V}(z) - K_H$ for $z \geq \beta$, and $\beta \to \underline{z}$.

[20] To apply Dini's Theorem, the function's domain must be compact. However, simply transform log-likelihood states, z, back into probability states, $p \in [0, 1]$, and, for all ϕ-values, extend the function to $p = 0, 1$ to preserve continuity.

For (iv), the pointwise convergence above $\underline{z}$ is immediate. For $z \leq \underline{z}$,

$$\begin{aligned} 0 \leq F_L(z) - V_L &= C_1^*(u_1 - 1)e^{u_1 z} \\ &= (K_H - V_L)\left(\frac{u_1 - 1}{u_1}\right)^{u_1}\left(\frac{V_H - K_H}{K_H - V_L}e^z\right)^{u_1} \\ &\to (K_H - V_L)e^{-1}\lim_{u_1\to\infty} y^{u_1}. \end{aligned}$$

The remainder of (iv) follows from $z < \underline{z} \Rightarrow y < 1 \Rightarrow y^{u_1} \to 0$ and $z = \underline{z} \Rightarrow y = 1 \Rightarrow y^{u_1} \to 1$. Finally, (v) is immediately implied by (iii) and (iv). □

Proof of Theorem 2. In the proposed equilibrium candidate, for all $z \in \mathbb{R}$, trade is immediate, $W(z) = F_L(z) = K_H$, and $F_B(z) = \bar{V}(z) - K_H$. As in the proof of Theorem 1, Eq. Conditions 1, 3, and 4 are by construction of the Σ-profile. In the candidate, $\beta = -\infty$, so verification of *Seller Optimality* (Eq. Condition 2) is trivial: for all z, $W(z) \leq K_H$, so for $\theta \in \{L, H\}$:

$$\sup_{\tau\in\mathcal{T}} E^\theta\left[e^{-r\tau}(W(Z_\tau) - K_\theta)\right] \leq K_H - K_\theta = F_\theta(z).$$

Finally, the verification of *Buyer Optimality* (Eq. Condition 5) is identical to the one given for the case of $z > \beta^*$ in the proof of Theorem 1.

To see that no other Σ-equilibrium exists, suppose first that $\Sigma(\beta, \dot{q})$ was an equilibrium with $\beta \in \mathbb{R}$. The analysis from Section 3.1 again applies, and therefore F_B, β, C_1, C_2 must satisfy (15)-(18). Solving the system, as in Lemma 1, gives the unique solutions as

$$\beta = \ln\left(\frac{K_H - V_L}{V_H - K_H}\right) + \ln\left(\frac{u_1}{u_1 - 1}\right),$$

which is not in $\mathbb{R}$ when the SLC fails (i.e., $K_H - V_L < 0$), contradicting the supposition. Finally, if $\beta = \infty$, then $F_B(z) = 0$ for all $z \in \mathbb{R}$. But then the buyer would improve his payoff by offering K_H (leading to payoff $\bar{V}(z) - K_H > 0$) for any z. Hence, no other Σ-equilibrium exists. □