WARWICK

ECONOMICS

CRETA

Centre for Research in Economic Theory and its Applications

Discussion Paper Series

Persuasion and Pricing : Dynamic Trading with Hard Evidence

Peter Eso & Chris Wallace

August 2016

No: 24

CRETA
Centre for Research in Economic Theory and its Applications

Department of Economics
University of Warwick, Coventry,
CV4 7AL, United Kingdom

Persuasion and Pricing:
Dynamic Trading with Hard Evidence[†]

Péter Eső[‡]
University of Oxford
peter.eso@economics.ox.ac.uk

Chris Wallace
University of Leicester
cw255@leicester.ac.uk

August 3, 2016

Abstract. In a simple trading game the buyer and seller have repeated opportunities to acquire and disclose hard (verifiable) evidence about the value of the tradable good. The parties disclose individually favourable information but conceal signals which are beneficial to the other side. In a leading case of interest with a finite horizon and sufficiently patient players, the equilibrium is characterized by a period of skimming (in which the seller makes offers acceptable only to informed buyers) concluded by a single settling period in which agreement is reached for sure. The length of delay until agreement and the corresponding efficiency loss are decreasing in the time horizon and in the abilities of the trading parties to identify the good's value, but increasing in impatience. An arbitrarily long time horizon can generate either immediate agreement or no trade between uninformed parties.

JEL Classification Codes. C78, D82, D83. **Keywords.** Persuasion, trading games, bargaining, hard evidence, information disclosure, skimming, settling, delay.

1. INTRODUCTION

Most game-theoretic models of bargaining abstract away from the role of *persuasion* (the production and disclosure of verifiable evidence) and instead focus on modelling the parties' bargaining power (as captured by the order and frequency of offers), risk attitudes, impatience, and unverifiable private information.[1] The goal of this paper is to capture within a trading game the players' *skill of discovering and proving facts* concerning the value or fair distribution of the surplus, and analyse its effects on bargaining outcomes.

[†]**Acknowledgements.** We thank seminar participants at Oxford and Leicester amongst (many) others for their questions and comments. Special thanks to Simon Board, Vincenzo Denicolò, Johannes Hörner, María Sáez Martí, David P. Myatt, Thomas Noe, Joel Sobel, and Eyal Winter for their helpful comments and advice.
[‡]**Corresponding Author.** Department of Economics, University of Oxford, Oxford OX1 3UQ, UK.
[1]For a textbook treatment of bargaining theory, see Muthoo (1999). Kennan and Wilson (1993) and Ausubel, Cramton, and Deneckere (2002) provide surveys of bargaining games with incomplete information.

In Section 2 a dynamic trading model is introduced in which, over time, both players may be able to generate *verifiable* but *concealable* evidence about the value of the good, while the seller is making repeated take-it-or-leave-it offers to the buyer. The equilibrium is characterized in Section 3. It determines the time and price of *settlement* (agreement in the absence of either side proving the good's value to the other), as well as the prices set by the seller before and, off the equilibrium path, after the time of settlement. The fundamental tension is that both parties correctly suspect the other of concealing disadvantageous evidence from the first moment they meet. However, they are also aware that any delay reduces the surplus that they might share because of discounting, or in the case of a finite horizon, a deadline.

In the leading case of interest analysed in Section 4, with a finite time horizon and sufficiently patient players, equilibrium play is characterized by a *finite delay* before settlement. The rough intuition is as follows. At the first feasible trading date, by assumption, both parties may already be concealing information beneficial for the other side. Hence, owing to adverse selection, a genuinely uninformed but 'suspicious' seller is unable to make an offer that a similarly uninformed and suspicious buyer would be happy to accept. However, if no settlement is ever reached, then (even with very little impatience) the seller would prefer to offer a discrete price cut in order to entice trade. In equilibrium the discount is offered gradually over time (in the form of delay, unless either party discovers that the good is valuable), until an endogenous settlement date is reached.

In the finite-horizon case the above-mentioned finite delay is *shortened*, and the corresponding efficiency loss decreased, by a greater ability of either party to identify the good's value. Perhaps somewhat surprisingly the delay is also shortened by a *longer* time horizon or a lesser degree of impatience. This is because any of these exogenous changes increases the seller's continuation payoff from 'skimming' (setting a high price that only a buyer privately informed of the good's high value is willing to pay). At the time of settlement the price must compensate the seller for this alternative continuation payoff. However, a higher settlement price requires earlier settlement because the uninformed and increasingly suspicious buyer's willingness to pay is falling over time.

In contrast to the positive but limited delay in the finite-horizon case, under an arbitrarily long (and approximating infinite) time horizon settlement may occur *immediately* or *never*. Section 5 demonstrates that the former occurs with arbitrarily frequent offers and any positive degree of impatience, whereas the latter occurs when the parties are infinitely patient. If the trading partners do not discount future payoffs then they may as well wait until the true value of the good is revealed and so trade can take place at a 'fair' price, hence they never settle. If, however, the players are even slightly impatient then the seller is willing to agree to a price cut, but the discount he is willing to offer tends to zero as the time horizon grows unboundedly.

The model is motivated by the importance of verifiable (perhaps documentary) evidence to practical bargaining situations such as those found in labour market negotiations or settlement talks in commercial disputes. Written accounts of (and advice for) real-life or imaginary negotiations often focus on the parties' ability to generate and disclose persuasive, hard information.[2] Nevertheless, bargaining theory to date has been largely silent on the potential effects of the presence of such evidence (a review of the literature is contained in Section 6). The key contribution of this paper therefore is to provide a formal model of such an environment, yielding testable comparative-statics results. For instance, the analysis here shows that more experienced bargainers (those who are more likely to be able to generate verifiable evidence) settle earlier; so do more patient parties, or those with a longer window of opportunity in which to conclude their negotiations.

Before proceeding with the general model and formal analysis, the remainder of the introduction provides a review of the basic idea. It does so in a relatively straightforward setting (where, among other simplifications, only the price-setting party may acquire information) that nevertheless captures some of the key mechanisms and intuitions.

The Basic Idea. Suppose that a seller (he) and a buyer (she), both risk neutral, have a limited time horizon (say, T periods) to trade a single, indivisible good. The monetary value of the good (and hence the surplus from trade) is determined by a binary state of nature: in the high state the buyer's valuation is 1 and the seller's is $v \in (0,1)$, whereas in the low state both valuations are zero.[3] The state is initially unknown to both players; the common prior that it is high is σ.

In each period the seller has a chance to learn the value of the good perfectly, and does so with a constant (and perhaps small) probability. The seller's information comes in the form of *hard evidence* which he can credibly disclose to the buyer. Equally, the seller could conceal any evidence which has been received; he cannot credibly communicate the absence of such evidence, however. As a result, the buyer is unable to tell whether the seller remains silent because he is uninformed, or because he chose to conceal the evidence. Following this, the seller offers the buyer a price which she may accept (ending the game) or reject (resulting in a further opportunity for the seller to observe the value,

---

[2] A fictional story with this feature, *Parson's Pleasure* by Roald Dahl first published in 1958, and appearing in Dahl (1979), is described at the end of the literature review in Section 6.

[3] That is, there is positive ex-post surplus from trading (for both parties, at an appropriate positive price) in the high state; there is neither a gain nor a loss (at zero price) in the low state. v can be interpreted as the salvage value for the seller (in the high state, when the good is valuable). This is less than 1 because the seller must seek a new buyer.

disclose any evidence held, and offer another price, up until T). If the good remains unsold at T then the seller realizes its value: v in the high and zero in the low state.

If, at any time, the seller discloses his information then the buyer's valuation for the good becomes commonly known. In the focal, seller-optimal equilibrium (with or without time preference), after such disclosure, the seller sets the price to 1 in the high state and 0 in the low state, and trade occurs immediately. Therefore, it is in the seller's interest to immediately disclose 'favourable' evidence that proves that the state is high. In contrast, it is weakly dominated for him to disclose that the good's value is zero.

With this information-disclosure policy in force, when the seller remains silent, the buyer must update her beliefs about the value of the good. Let β_t denote her belief that the state is high at time $t \leq T$, conditional on observing no disclosure at or before t. The seller's silence may stem from his lack of information or his concealment of evidence that the state is low. Thus, over time, the buyer's belief must deteriorate: β_t is decreasing in t.

In the final period an uninformed seller has no further opportunity to obtain information about the value of the good. His expected salvage value from keeping the good is therefore σv. Given that β_t converges to zero, for a sufficiently long time horizon, it will be the case that $\sigma v > \beta_T$. In other words, an uninformed seller would not wish to offer a price that the buyer would be willing to accept at T. Working backwards (and, for simplicity, assuming no discounting), in the period immediately before the end an uninformed seller must have a higher continuation value: there is still an opportunity to discover the state is high and thereby extract maximal surplus from the buyer. In general, the continuation value of an uninformed seller planning to sell (at price 1) only if discovering favourable evidence is decreasing over time, as there is a diminishing number of opportunities left.

Before the game starts, a seller who expects to sell the good only when able to disclose favourable evidence must have a continuation value below σ: there is always the possibility that even though the state is high, he never discovers it and ends up keeping the good which is worth $v < 1$ to him. In sum, the uninformed seller's continuation value from selling only following a favourable disclosure is decreasing and below σ everywhere. In contrast, the buyer's belief β_t starts at σ (her prior) and ends up below σv, provided the horizon is sufficiently long (assumed in the rest of the section).

It follows that there must be a last moment (call it m) at which the seller's continuation value from selling only after disclosing favourable evidence is below the buyer's willingness to pay. After m, the seller's continuation value from selling only when discovering that the good's value is high is greater than the buyer's willingness to pay: from this point on the seller will sell only when discovering that the good's value is high. At m, however, the uninformed seller could offer a price exactly equal to β_m, the highest price the buyer will accept at m. Indeed, the seller would now do so: his continuation value from continuing to delay is exceeded by the buyer's willingness to pay.

Working backwards from m, the seller's continuation value at $t < m$ is somewhat different, because he anticipates selling the good for β_m with certainty at m. At some t prior to m, the seller who plans to wait till m to sell for sure (at price β_m), and before then sells (at price 1) only when discovering favourable evidence has continuation value

$$V_t^m = \sigma X(1 - \beta_m) + \beta_m,$$

where X stands for the probability that the seller discovers the value of the good between times t and m. That is, the seller receives a premium of $(1 - \beta_m)$ with probability σX on top of the price he will receive for sure (β_m) at m. On the other hand, at such a t, the buyer is prepared to pay only

$$\beta_t X(1 - \beta_m) + \beta_m,$$

which includes a premium over β_m (reflecting the fact that the buyer knows her surplus will be extracted if the state is high and the seller discovers it between t and m). The price the buyer is willing to pay is strictly less than V_t^m, since $\beta_t < \sigma$ for all $0 < t < m$.[4]

This argument implies that the seller will wait to 'settle' for sure in period m, at price β_m, trying in the meantime to generate hard evidence of the good's high value. Off the equilibrium path (after m) the seller would have continued to wait for hard evidence, selling at a high price only upon its discovery, eventually keeping the good and obtaining its salvage value at T when no evidence was forthcoming.

This "wait-settle-wait" structure is present in the equilibrium of the more general model analysed in this paper. So too are the comparative-static intuitions that accompany it. Consider, for instance, extending the time horizon. Allowing the players more time to trade does not affect the evolution of the buyer's belief over time, β_t. It does, however, increase the number of chances the seller gets to discover that the state (hence the good's value and the trading surplus) is high. This raises his continuation value (from sale only after disclosure) at any t after m, meaning the buyer must pay a higher price. Necessarily this must occur earlier, since the buyer's willingness to pay is falling: so delay is reduced.

Similarly, a seller with a higher valuation for the good in the high state (a higher salvage value) also brings forward the date of (certain) trade.

[4] In this basic setting the players are infinitely patient, and the buyer has no opportunity to discover the good's value for herself. Both these restrictions are lifted in the general model. Now there is also a risk for the seller that the buyer discovers the good has no value: this affects the premium accruing to the seller in his continuation value from waiting, as does impatience; and the comparison above becomes more subtle.

Finally, a seller who is better able to identify the quality of the good (one who has a higher probability of generating hard evidence) will also settle earlier: not only is the seller's continuation value raised from the increased chance of observing the good's quality in any period, but also the buyer's beliefs fall more rapidly (since she becomes more pessimistic even more quickly upon observing no disclosure). In order to obtain the higher price that the seller can now guarantee he must bring forward the sale even further, before the buyer's willingness to pay drops too far.

These observations and intuitions crop up again in the analysis of the general model (to which the paper now turns), albeit with some modification and qualification. One key difference is that the above argument abstracts from the buyer's ability to obtain hard information: this has an important effect upon the wait-settle-wait structure described above. Now the buyer herself may be concealing evidence, and it will be in the seller's interest to attempt to sell to this type of buyer when waiting to settle. Moreover the buyer's ability to identify the state of nature will have an important influence on the timing of agreement, the price of trade, and the surplus the parties are able to extract. In the unbounded horizon case (again unmentioned above), the two parties may be either unable to settle at all or might wish to settle right away, depending upon their time preferences.

2. A Trading Game with Hard Evidence

A seller (S, he) has an indivisible good for sale to a single buyer (B, she). The state of the world is $\omega \in \{H, L\}$. If $\omega = H$ then the good's monetary value is 1 for B and $v \in (0, 1)$ for S; if $\omega = L$ then it is 0 for both players. v is commonly known, whereas the realization of ω is initially unknown to both parties; they have a common prior, $\sigma = \Pr[\omega = H] \in (0, 1)$.

2.1. Information Structure and Timing of Actions.

Publicly-observable physical time is indexed by $t \in \{\Delta, 2\Delta, \ldots\}$. Discrete time makes the order of moves clear and the players' strategies well-defined. $\Delta > 0$ is treated as *small* throughout; the case of *arbitrarily frequent interaction* (discrete time with $\Delta \to 0$) will be studied for being inherently interesting and eminently amenable to analysis.

At the beginning of each period (or equivalently, during the preceding Δ amount of time) S learns the state of the world with probability $r_S\Delta$, and otherwise remains uninformed. Independently, B learns the state with probability $r_B\Delta$ and otherwise remains uninformed. A player may either *verifiably disclose* or *conceal* knowing the state; however, a *lack of information* about the state *cannot be proved*.

Formally, at the beginning of each period t, S observes $s_t \in \{L, H, \emptyset\}$ and B observes $b_t \in \{L, H, \emptyset\}$, such that

$$\Pr[s_t = \omega] = 1 - \Pr[s_t = \emptyset] = r_S\Delta \quad \text{and} \quad \Pr[b_t = \omega] = 1 - \Pr[b_t = \emptyset] = r_B\Delta.$$

$s_t, b_t \in \{H, L\}$ are verifiable, but $s_t = \emptyset$ and $b_t = \emptyset$ are not. When a nonempty s_t or b_t is disclosed, ω becomes commonly known to the players.[5] The idea is that S (respectively, B) has the ability to generate irrefutable evidence regarding the good's value at a rate r_S (respectively, r_B) per unit time, but when a player finds such evidence (for instance, an original document demonstrating the good's provenance) he or she is able to conceal it.

Assume that S and B may only trade during the time interval $[\underline{t}, \bar{t}] \subseteq [0, \infty]$, which is exogenously given; let $\mathbb{T} \equiv [\underline{t}, \bar{t}] \cap \{\Delta, 2\Delta, \ldots\}$. At each $t \in \mathbb{T}$ the players may first simultaneously disclose any evidence obtained at or before t, then S makes a take-it-or-leave-it price offer, p_t. If B accepts then she pays p_t and realizes the good's value; if she rejects then play continues in the next period. If the offer is rejected in the final period then S keeps the good. Both players are risk neutral with respect to monetary gains, and discount payoffs at a rate $\rho \geq 0$ per unit time.

Note that S and B may obtain information before they meet and start trading: even if $\underline{t} = 0$, the first trading period is $t = \Delta$, and so when S sets the price, the players may already be concealing a signal received in that period. This can create a "wedge" between the players' initial estimates regarding the good's value: if S is believed to conceal "bad" signals and B to hide "good" ones (which is indeed in each player's self interest, as S wants to sell high and B to buy low), then a genuinely uninformed S will be more optimistic about $\omega = H$ than a genuinely uninformed B when the first offer is made. Indeed, the larger $\underline{t}$ is, the larger is the wedge between the players' initial beliefs. (All this is, of course, conceptually different from commonly known but different priors.)

2.2. **Solution Concept.** Equilibrium henceforth refers to pure-strategy perfect Bayesian equilibrium as defined in Fudenberg and Tirole (1991, Definition 8.2), suitably extended to the present context with a continuum of actions. Throughout, focus will be on the (uninformed) seller-optimal equilibrium of this class. In addition, the following two refinements on out-of-equilibrium beliefs will be imposed. First, if S makes an out-of-equilibrium price offer, then B's beliefs are required to satisfy the *Intuitive Criterion*. (That is, B must put zero weight on a type of S that could not possibly gain from the deviation for any sequentially rational response of B.) Second, if B strays from equilibrium by

[5]Note players cannot obtain contradictory signals (e.g., $s_t = H$ and $b_{t'} = L$), and this is commonly known. The assumption that signals (when obtained) are perfect is made for analytic tractability: considering similar models with imperfect hard evidence would be an interesting topic for future research.

rejecting a price that she should have accepted for sure, irrespective of her private information, then S infers that B is uninformed. As will be clear from the analysis, this is a sensible and intuitively appealing assumption.[6]

If S observes $s_t = H$ then it is optimal for him to disclose this fact and charge $p_t = 1$, which ought to be accepted by B without delay.[7] We shall make the analogous assumption that if B finds $b_t = L$ then she discloses it inducing S to set $p_t = 0$ for immediate trade as well. This can be thought of as an *indifference-breaking condition* as both parties earn zero surplus whether or not they trade at zero price in state $\omega = L$. (An alternative assumption that would work equally well is that B "walks away" from the relationship immediately once she learns $\omega = L$, because she has other, more productive things to do.)

As a result of the observation and indifference-breaking assumption made in the previous paragraph, at any $t \in \mathbb{T}$ and absent prior disclosure, there may be two types only of B present: the uninformed buyer, labelled $B_\emptyset$, and an informed buyer who knows the state is high, labelled B_H. Equally, there may be two types only of S present: an uninformed $S_\emptyset$ and an informed seller who knows the state is low, S_L.

The fact that *self-serving* signals are immediately disclosed (to end the game) whilst *self-defeating* ones are concealed (to keep the other player guessing) are intended, key features of our model of persuasion and pricing. The seller's objective is to find an optimal price-discrimination strategy knowing that he faces a buyer who is either uninformed or concealing an H-signal (if she is informed), whilst knowing that an uninformed B suspects him of potentially concealing an L-signal. Since S knows that at any $t \in \mathbb{T}$ he faces either an uninformed buyer ($B_\emptyset$) or one who knows the good is valuable (B_H), he must either charge a relatively high price that only the latter type accepts (called *skimming*), or a lower price that both buyer types accept for sure (called *settling*).

In the next section we construct the unique equilibrium satisfying the above requirements.

3. Equilibrium Construction

In this section an equilibrium is characterized in which at each $t \in \mathbb{T}$ the parties disclose all *self-serving* signals received at or before t and, conditional on such disclosure, immediately trade at a price equal to B's valuation. If no such persuasive disclosure is made then both types $S_\emptyset$ and S_L of the seller either *skim* (offer a price p_t acceptable only to B_H),

[6]In particular, if S makes an offer that an uninformed buyer is just indifferent between accepting and rejecting then a buyer who has previously observed $\omega = H$ has a strict incentive to accept.

[7]Immediate disclosure of $s_t = H$ is *strictly* optimal for S only if $\rho > 0$, but for continuity it is assumed if $\rho = 0$ as well. The fact that after such disclosure B accepts $p_t = 1$ right away follows from the usual reasoning that otherwise S could slightly lower the price to ensure acceptance.

or *settle* (offer a price q_t acceptable to both $B_\emptyset$ and B_H). In case, out of equilibrium, a settlement price is rejected, then S believes B is uninformed and makes no inference regarding the state. In addition, after an out-of-equilibrium price proposal B's beliefs satisfy the Intuitive Criterion. *Self-defeating* signals are concealed in equilibrium.[8]

Proposition 1 establishes the existence and uniqueness of an equilibrium with these properties. The equilibrium either prescribes skimming in all periods, or (e.g., if the horizon is sufficiently long but finite) involves a *final period of settlement* m, after which, off the equilibrium path, S would skim in all periods. Before reaching m, in the absence of signal disclosure by either party, the seller skims in all periods provided the time horizon is finite and the players are sufficiently patient.

Both on and off equilibrium path, S_L pools with $S_\emptyset$ when offering skimming and settling prices. This is not assumed: there is no equilibrium in which S_L separates from $S_\emptyset$.[9]

3.1. **Equilibrium Characterization.** In order to derive m (the last settlement date) two quantities are defined: an uninformed buyer's willingness to pay, β_t, and an uninformed seller's expected payoff from skimming at t and in all subsequent periods, V_t.

The good's expected value for a genuinely uninformed B (absent any disclosure by S at or before t) may be computed using Bayes' rule, for all $t \in \{\Delta, 2\Delta, \ldots\}$,

$$\beta_t = \Pr[\omega = H \mid B_\emptyset, \forall \tau \le t : s_\tau \in \{\emptyset, L\}] = \frac{\sigma(1 - r_S\Delta)^{t/\Delta}}{\sigma(1 - r_S\Delta)^{t/\Delta} + 1 - \sigma}. \tag{1}$$

β_t is strictly decreasing in t: an uninformed buyer loses confidence in the good's value as S remains quiet (he does not prove $\omega = H$); for B it becomes increasingly likely that S is in fact concealing a low signal. As $t \to \infty$, this expected valuation tends to zero. Of course, as time passes, it also becomes less likely that B remains uninformed.

The seller believes that his price proposal is only ever rejected by buyer-type $B_\emptyset$ (on or off the equilibrium path, S having skimmed or attempted to settle), hence right before the beginning of period t, i.e., before s_t and b_t are observed, he believes B is uninformed. Therefore, if the seller remains uninformed at t and the buyer makes no disclosure either, then he ($S_\emptyset$ in period t) believes the state is $\omega = H$ with probability

$$\hat{\sigma} = \Pr[\omega = H \mid S_\emptyset, b_t \in \{\emptyset, H\}, \forall \tau < t : b_\tau = \emptyset] = \frac{\sigma}{\sigma + (1 - \sigma)(1 - r_B\Delta)}.$$

[8]As will be seen, since disclosing $b_t = H$ results in immediate trade at price 1, B is strictly worse off (before m: see the next paragraph) or indifferent (after m). An analogous argument applies for S observing $s_t = L$.

[9]If S_L separated from $S_\emptyset$ in equilibrium at $t > m$ then $B_\emptyset$ would accept a settlement price $q_t = \sigma$ from $S_\emptyset$; this would be preferred to skimming by $S_\emptyset$ for $\Delta > 0$ sufficiently small hence unraveling the equilibrium. On the other hand, prior to any such m, S_L would have an incentive to deviate from separation to mimic $S_\emptyset$ and obtain a positive payoff with some probability (i.e. when m is reached).

Note that $\hat{\sigma}$ is a constant (the same for all t) that exceeds σ but converges to it as $\Delta \to 0$. Intuitively, $\hat{\sigma} > \sigma$ as B might have learned $\omega = H$ exactly at t, but the chance of that happening is infinitesimal as the period length goes to zero.

In the final trading period, at $T = \max \mathbb{T}$ (which coincides with $\bar{t}$ unless $\bar{t}$ is not divisible by Δ) the seller may skim using $p_T = 1$ as B_H has zero continuation value from rejecting the final offer. The expected payoff of $S_\emptyset$ from skimming at T is

$$V_T = \hat{\sigma}\left[r_B\Delta + (1 - r_B\Delta)v\right], \tag{2}$$

because, conditional on $\omega = H$ (i.e., with probability $\hat{\sigma}$) $S_\emptyset$ sells to B_H at price 1 with probability $r_B\Delta$ otherwise he keeps the good that is worth to him v; whereas if $\omega = L$ then the seller gets no payoff from skimming in the last period. For any $t < T$, if the seller plans to skim with price 1 in all subsequent periods then he can skim at t using price $p_t = 1$ as well, and his expected payoff from doing so is V_t satisfying

$$V_t = \hat{\sigma}\left[r_B\Delta + \frac{(1 - r_B\Delta)r_S\Delta}{1 + \rho\Delta}\right] + \frac{(1 - r_B\Delta)(1 - r_S\Delta)}{1 + \rho\Delta}V_{t+\Delta}. \tag{3}$$

If the state is high (which has probability $\hat{\sigma}$ from the perspective of $S_\emptyset$) then B accepts the skimming price with probability $r_B\Delta$, or else, if she is not of type B_H already, then $S_\emptyset$ may become informed of $\omega = H$ within the next period and sell the good at price 1 with a one-period delay. If either player learns $\omega = L$ within the next period then the seller's payoff is zero. If neither player learns the state by next period then $S_\emptyset$ will get continuation payoff $V_{t+\Delta}$ with Δ delay. The expressions in (2) and (3) uniquely define V_t.

If $S_\emptyset$ skims at t using $p_t = 1$ (anticipating that he will skim, absent disclosure, in all future periods as well) then his expected payoff is V_t; alternatively he can settle with price (and sure payoff) $q_t = \beta_t$. Hence, provided he will skim in all future periods, $S_\emptyset$ prefers skimming to settling at t if and only if $V_t > \beta_t$. This condition holds in the final period if the time horizon is sufficiently long because V_T is a positive constant (irrespective of T) whereas β_t converges to zero as $t \to \infty$.

If $V_t > \beta_t$ for all $t \in \mathbb{T}$ then let $m = 0$, otherwise define

$$m = \max\left\{t \in \mathbb{T} \,|\, V_t \leq \beta_t\right\}. \tag{4}$$

This is the *final moment of settlement* (the last time of certain trade without disclosure). By construction $V_t > \beta_t$ holds in all periods $t > m$, hence $S_\emptyset$ skims using $p_t = 1$ provided neither player reveals the state at or before t. Note that S_L also skims at all $t > m$: even though he expects zero payoff by doing so, a deviation (from which only S_L could gain) would reveal his type by the Intuitive Criterion and result in zero payoff as well.

In periods before m the comparison between V_t and β_t is no longer relevant because S is not expected to skim using price 1 in all subsequent periods. In fact, S is expected to settle in m at the latest. If settlement in the absence of disclosure is anticipated after t then B's continuation value from rejecting the offer at t is positive, which lowers the maximal price that she accepts at t.

Indeed, suppose that in period t it is anticipated (along the equilibrium path) that S will next settle in period $k > t$, at price q_k, and he will skim for all $\tau \in (t,k) \cap \mathbb{T}$. Then the highest price B_H accepts at t (i.e., the skimming price at t) is

$$p_t = 1 - \left(\frac{1 - r_S\Delta}{1+\rho\Delta}\right)^{(k-t)/\Delta} (1 - q_k). \tag{5}$$

To see why, note that B_H expects a surplus of $(1 - q_k)$ at k (that is, in $(k-t)/\Delta$ periods) unless S also discovers that the state is high, in which case she gets nothing. B_H does not expect any additional surplus before k as she will be skimmed.

By similar reasoning, the highest price that $B_\emptyset$ accepts (i.e., the settlement price at t, anticipating the next settlement at $k > t$) is

$$q_t = \beta_t - \left(\frac{1 - r_S\Delta}{1+\rho\Delta}\right)^{(k-t)/\Delta} \beta_t(1 - q_k) + \left(\frac{1 - r_B\Delta}{1+\rho\Delta}\right)^{(k-t)/\Delta} (1 - \beta_t)q_k. \tag{6}$$

Again, the intuition is that $B_\emptyset$ has to be compensated for the continuation value she could realize by refusing to buy the good of expected value β_t at t. This surplus is $(1 - q_k)$ in period k if the state is high (i.e., with probability β_t) provided S does not discover the state by then, less q_k if the state is low (i.e., with probability $(1 - \beta_t)$) provided B herself does not discover the state by then. It is easy to see that $q_t < p_t$, that is, the skimming price (targeting B_H only) is indeed greater than the settlement price (targeting both types of B) when the next settlement is anticipated at price q_k in period k.

With the skimming and settlement prices in hand it is straightforward to write down the uninformed seller's expected payoff from skimming versus settling at t when the next settlement (in equilibrium, in the absence of disclosure) is anticipated in period k. The comparison of these two payoffs determines whether S skims or settles at t. Working backwards, the seller's optimal decision can be derived at $t - \Delta$, and so on, which yields the unique equilibrium satisfying the properties described in Section 2.

With the details of the rest of the derivation relegated to the Appendix, the following Proposition characterizes the desired equilibrium of the model.

Proposition 1 (Equilibrium Characterization). *There exists a unique equilibrium such that*

(i) *S conceals $s_t = L$ and discloses $s_t = H$, while B conceals $b_t = H$ and discloses $b_t = L$; trade occurs immediately after a hard signal disclosure at a price matching B's valuation;*

(ii) *B's out-of-equilibrium beliefs satisfy the Intuitive Criterion after S's deviation; whereas $S_\emptyset$ believes B is uninformed if she rejects a price she was supposed to have accepted for sure.*

For all $t \in \mathbb{T}$, absent prior disclosure the seller either skims (sets a price p_t accepted by B if and only if $b_\tau = H$ for some $\tau \leq t$), or settles (sets a price q_t accepted by B after any history).

Define $m = 0$ if $V_t > \beta_t$ for all $t \in \mathbb{T}$, otherwise let $m = \max\{t \in \mathbb{T} \,|\, V_t \leq \beta_t\}$. In the absence of prior disclosure by either player, S skims with price $p_t = 1$ at any $t > m$. If $m \in \mathbb{T}$ then S settles with price $q_m = \beta_m$ at $t = m$. For ρ sufficiently small the seller skims in all periods $t < m$.

The main difference between the structure of the equilibrium described in the basic idea of Section 1 and the one of the above proposition is that, rather than *delay*, an uninformed seller now *skims*, hoping to sell early to an informed buyer concealing $\omega = H$. Otherwise, for ρ small enough, the equilibrium retains the wait-settle-wait structure outlined earlier. The possibility that the buyer can discover the state of nature has other effects as well, discussed later alongside the comparative statics of Section 4. For now however, notice that the condition that players be sufficiently patient is natural: in the periods directly before m, the seller is willing to skim (rather than settle immediately) only if the probability that he (or, now, the buyer as well) discovers $\omega = H$ in the intervening time outweighs his impatience for waiting until m. The risk that the buyer observes $\omega = L$ is, of course, reflected in both the seller's continuation value of waiting till m and the highest price the buyer is willing to pay, leaving the comparison between the two unaffected.

3.2. **Arbitrarily Frequent Interactions.** In the next two sections equilibrium play and the corresponding quantities of interest are further analysed in two, leading cases of interest: (i) with a *finite time horizon* and *sufficiently patient players* (in which case S skims either to the end or until a unique, final moment of settlement), and (ii) with an *arbitrarily long time horizon*, in which case (as shall be seen) S skims in all periods if $\rho = 0$, but settles right away for any $\rho > 0$ if the period length ($\Delta > 0$) is sufficiently short. In both cases equilibrium price paths and comparative statics results are derived and illustrated.

These cases will be analysed under the assumption of *frequent interaction*. The model where $\Delta > 0$ is arbitrarily small is not only theoretically attractive (as a valid approximation of continuous time) but also results in formulæ that make the calculations more palatable. In practical terms, formulæ for price paths and various constraints will be evaluated in the limit as $\Delta \to 0$. It should be noted that our model is still a properly specified, discrete time game, with the equilibrium characterized by Proposition 1 for any $\Delta > 0$.

By convention, *subscript* t will be used to refer to the value of a variable in period t of the discrete-time model with a given $\Delta > 0$; however, a variable written as a function with *argument* t will denote the limit of the variable's value at t as $\Delta \to 0$. For example,[10]

$$\beta(t) \equiv \lim_{\Delta \to 0} \beta_t = \frac{\sigma}{\sigma + (1-\sigma)e^{r_S t}}. \tag{7}$$

It can be shown (see the Appendix) that as $\Delta \to 0$, V_t defined by (2)-(3) converges to

$$V(t) = \left[1 - e^{-(r_B + r_S + \rho)(\bar{t} - t)}\right] \sigma\lambda + e^{-(r_B + r_S + \rho)(\bar{t} - t)} \sigma v, \tag{8}$$

where $\lambda = (r_B + r_S)/(r_B + r_S + \rho) \in (0, 1]$. $V(t)$ is decreasing and concave for ρ sufficiently low. Note that if $\bar{t} = \infty$ (i.e., if the time horizon is unbounded), then the expected payoff to $S_\emptyset$ from skimming forever becomes constant in t, $V(t) = \sigma\lambda$.

The final moment of settlement, m as defined by (4) provided it exists, converges as $\Delta \to 0$ to the solution of the equation

$$m^* = \bar{t} - \frac{1}{r_B + r_S + \rho} \ln\left(\frac{\lambda - v}{\lambda - \beta(m^*)/\sigma}\right). \tag{9}$$

Note that $V(0) < \sigma = \beta(0)$, whereas $V(\bar{t}\,) > \beta(\bar{t}\,)$ for $\bar{t}$ sufficiently large, in particular,

$$\bar{t} > \frac{1}{r_S} \ln\left(\frac{1 - \sigma v}{(1-\sigma)v}\right), \tag{10}$$

hence $m^* \in (\underline{t}, \bar{t}\,)$ exists if (10) holds and $\underline{t}$ is sufficiently close to 0.

By Proposition 1, S skims for all $t < m$ provided ρ is sufficiently low. The $\Delta \to 0$ limit of the pre-settlement skimming price in (5) for $k = m^*$ is

$$p(t) = 1 - [1 - \beta(m^*)]\, e^{-(r_S + \rho)(m^* - t)}. \tag{11}$$

Note that $p(t)$ is decreasing in t. The skimming price charged by an uninformed seller falls over time because the informed buyer at whom the price is targeted becomes more and more willing to wait until m^* and take advantage of the settlement price she will face then. The buyer knows the seller has fewer and fewer opportunities to discover $\omega = H$. Therefore, the seller has to offer a larger and larger discount in his skimming price in order to ensure acceptance by the buyer. Naturally, the price the informed buyer is willing to pay must converge to $\beta(m^*)$ at $t = m^*$, which is confirmed by inspection of (11).

[10]This and all the formulæ for $\Delta \to 0$ below are derived in the Appendix.

4. Waiting to Settle

Recall that if the time horizon is sufficiently long then a final moment of settlement m exists; moreover, by Proposition 1, if the players are sufficiently patient (i.e., $\rho \leq \bar{\rho}$ for some $\bar{\rho} > 0$) then S skims in all trading periods before m. That is, in equilibrium (absent signals disclosure and hence immediate trade) the players are *waiting to settle* for a specific amount of time, m.

Figure 1 illustrates this case with a numerical example. The uninformed buyer's belief is drawn in red; the uninformed seller's continuation value from skimming until $\bar{t}$ in blue.[11] m^* is the last moment where the former is above the latter. Thus, the seller skims (out of equilibrium) from m^* to $\bar{t}$. At m^*, the settlement price is $\beta(m^*)$; the pre-settlement skimming price given by $p(t)$ is drawn in black.

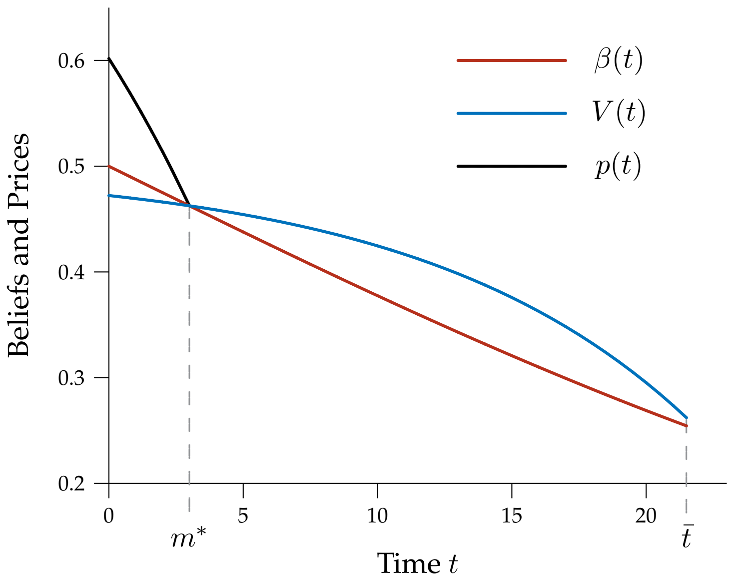

FIGURE 1. Waiting to Settle with a Finite Horizon.

The figure illustrates a case of waiting to settle with a finite horizon $\bar{t} = 22$. The parameter values are $r_B = r_S = 0.05$, $v = \sigma = 0.5$, and $\rho = 0$. The period length $\Delta > 0$ is taken to be arbitrarily short. With these data $\beta(t) < V(t)$ for all $t > m^* \approx 3$. So S skims (absent disclosure) with price $p(t)$ for all $t < m^*$, settles at m^* with price $\beta(m^*) \approx 0.46$, and (off equilibrium path) skims for all $t > m^*$.

The next result establishes comparative statics of m^*, $\beta(m^*)$ and $p(t)$ for all $t < m^*$ in the parameters of the model. Formal proofs for the case of arbitrarily frequent interaction (that is, in the limit as $\Delta \to 0$) are given in the Appendix.

[11] $V(t)$ is drawn in blue for all $t \in [0, \bar{t}]$, whereas of course it is a "valid" continuation value only for $t \geq m^*$.

Proposition 2 (Finite Horizon: Comparative Statics)**.** *Assume $\bar{t} < \infty$ is sufficiently high, $\rho \geq 0$ is sufficiently low so that a final moment of settlement exists, and consider arbitrarily frequent interactions ($\Delta \to 0$).*

The $\Delta \to 0$ limit of the last moment of settlement, m^*, is increasing in the prior σ and discount rate ρ, but decreasing in the seller's outside option v, the time horizon $\bar{t}$, and the probabilities r_B and r_S that the players obtain hard evidence.

The $\Delta \to 0$ limit of the settlement price, $\beta(m^*)$, is increasing in v, $\bar{t}$, r_B and r_S and decreasing in ρ. In contrast, the skimming price before settlement, $p(t)$ from (11)*, uniformly decreases for all $t < m^*$ for an increase in v, $\bar{t}$, r_B and for a decrease in ρ.*

Some of these comparative statics effects can be established simply by "shifting curves" (in particular, just the graph of V) in Figure 1, hence these are valid no matter how large Δ is. For example, an increase in $\bar{t}$ shifts the graph of V to the right horizontally, without changing its shape, as $V(\bar{t}\,) \equiv \sigma v$ and $V(t)$ depends on $(\bar{t} - t)$ not t itself. The graph of β is unchanged for a change in $\bar{t}$, so the last moment of settlement m^* where V and β intersect must decrease. The settlement price goes up along the β curve as m^* falls; however, the skimming price at every $t < m^*$ *decreases* (because B's continuation value from rejecting this price goes up). An increase in v or a decrease in ρ also shifts the graph of V up (out) without affecting β, hence the equilibrium value of m^* goes down, $\beta(m^*)$ goes up, and $p(t)$ for $t < m^*$ goes down.

The results that a longer time horizon (more time to "haggle"), or a lower discount rate (more patience), or an increase in the seller's valuation for the good when it is valuable all bring the time of the settlement *earlier* (at a higher price) may seem counterintuitive. However, they all follow from the intuition that at the time of settlement the seller must be compensated for an increase in his expected payoff from skimming till the end (owing to higher $\bar{t}$ or v or lower ρ), and the only way to achieve that is by settling earlier, while the uninformed buyer's valuation is greater.

The fact that the time of agreement moves forward when r_B increases is equally straightforward: when the buyer becomes better able to identify the value of the good, the seller is unwilling to wait for settlement so long (his continuation value $V(t)$ is raised, while $\beta(t)$ is unaffected). The buyer faces lower skimming prices $p(t)$ before settlement: a buyer who is better able to identify ω can drive down the price she faces. However, the settlement price agreed will be higher as settlement takes place earlier while the buyer is still relatively optimistic about the good's value.

Likewise, an increase in r_S brings forward the date of settlement (and raises the associated price). Now $\beta(t)$ is affected as well as $V(t)$: the buyer's beliefs deteriorate faster because nondisclosure by the seller is more indicative of bad news. In order to charge a high price an uninformed seller needs to act relatively quickly. The fact that the seller's continuation value is also raised by an increase in his ability to identify ω reinforces this effect.

Finally, and more subtly, an increase in σ raises m^*. Although it certainly raises $V(t)$, as a quick inspection of (8) verifies, raising σ also flattens (and raises) $\beta(t)$. In particular, the buyer, being more sure of $\omega = H$, is relatively prepared to ignore silence from S and wait. The seller can therefore wait longer even though the rise in his continuation value entails a relatively high price to compensate for skimming after m^*. While $V(t)$ rises one-to-one with σ, and in parallel $\beta(t)$ rises with σ, the flattening of the buyer's beliefs comes in addition, thereby lengthening the time to agreement. Players who are more optimistic (yet uninformed) about the value of the good are prepared to spend longer waiting before reaching agreement, a fact which has implications for the induced level of inefficiency: a theme to which the paper now turns.

The next proposition concerns total surplus (or social welfare). If a settlement date exists and S skims leading up to it then the *delay* from skimming when the good is valuable (and provided $\rho > 0$) is the only the source of efficiency loss. At the start of period $t < m$ the expected total surplus conditional on $\omega = H$ but neither player knowing this fact is

$$w_t = [r_S\Delta + (1 - r_S\Delta)r_B\Delta] + (1 - r_S\Delta)(1 - r_B\Delta)\frac{w_{t+\Delta}}{1 + \rho\Delta}, \tag{12}$$

where the first (bracketed) expression is the probability that either player discovers $\omega = H$ at the beginning of period t and so B buys right away, whereas the second term is the continuation value of the total surplus in case both players remain uninformed. This difference equation and boundary condition $w_m = 1$ (from settling at m) determine w_t, the expected total surplus at t conditional on $\omega = H$ not known to either player.

For arbitrarily frequent interaction, i.e., in the limit as $\Delta \to 0$, the expected total surplus at t conditional on $\omega = H$ not being known to either player converges to

$$w(t) = \lambda + (1 - \lambda)e^{-(r_B + r_S + \rho)(m^* - t)},$$

with $\lambda = (r_B + r_S)/(r_B + r_S + \rho)$, as before. This can be easily verified by first computing

$$\dot{w}(t) = \lim_{\Delta \to 0} \frac{w_{t+\Delta} - w_t}{\Delta} = -(r_B + r_S) + (r_B + r_S + \rho)w(t)$$

from (12) and integrating it up with boundary condition $w(m^*) = 1$.

The ex-ante expected social surplus, $W \equiv \sigma w(0)$, is therefore

$$W = \sigma\left[\lambda + (1 - \lambda)e^{-(r_B + r_S + \rho)m^*}\right]. \tag{13}$$

The following Proposition concerning the comparative statics of W in various parameters is proved in the Appendix.

Proposition 3 (Finite Horizon: Welfare)**.** *Assume $\bar{t} < \infty$ and that a final settlement period m exists. Total surplus is increasing in the probabilities r_B and r_S that the players obtain hard evidence and decreasing in the discount rate ρ. It is increasing in the seller's outside option v and the time horizon $\bar{t}$. Increasing the prior σ can reduce W.*

Greater probabilities of the players obtaining hard evidence bring the settlement date earlier as well as increase the probability of trade from discovering $\omega = H$; total surplus increases on both accounts. Decreasing ρ brings the settlement earlier and decreases the cost of delay, too. Increases in v and $\bar{t}$ only affect m^* (reducing delay), hence both are socially beneficial. However, a higher σ leads to a longer delay, for reasons discussed above, which may outweigh the increase in the probability that the good is valuable.

5. INFINITE HORIZON

This section analyses the limit of equilibrium behaviour as the time horizon converges to infinity.[12] Limiting equilibrium behaviour is sensitive to certain parameter values: in particular, the equilibrium under "infinite horizon" exhibits *immediate settlement* for any $\rho > 0$ provided $\underline{t} = 0$ and $\Delta > 0$ is arbitrarily small; in contrast, the seller *skims forever* if $\rho = 0$. Of course, these features can be understood as rather natural limits of the equilibrium structure already described earlier in Proposition 1.

For a fixed period length $\Delta > 0$ the uninformed seller's expected payoff from skimming from t on forever, denoted by $\tilde{V}_t$ (where the "tilde" decoration refers to the infinite horizon case), has to be constant in t because the future looks the same at t and t' for $S_\emptyset$ after a history of no prior signal disclosure and B having rejected all previous prices. That is, $\tilde{V}_t$ is the solution to (3) with $\tilde{V}_t = V_t \equiv V_{t+\Delta}$. Therefore

$$\tilde{V}_t = \frac{r_B + r_S + (\rho - r_S) r_B \Delta}{r_B + r_S + \rho - r_B r_S \Delta} \hat{\sigma}$$

where, recall, $\hat{\sigma} = \sigma / [1 - (1 - \sigma) r_B \Delta] > \sigma$ is the uninformed seller's belief that $\omega = H$, which exceeds the prior because B might be concealing a just-received H signal.

A first observation is that when $\bar{t}$ grows unboundedly, with frequent interaction and sufficiently patient players, a final moment of settlement fails to exist. That is, in equilibrium S skims forever.[13] To see this note that for $\rho = 0$ the seller's payoff from skimming forever

[12] More precisely, the analysis here considers the case of an arbitrarily long time horizon with $\bar{t} \to \infty$.

[13] In contrast, recall that m^* solving (9) exists when $\bar{t}$ is finite but sufficiently large.

becomes $\tilde{V}_t = \hat{\sigma}$, which *strictly* exceeds β_t for all $t = \{\Delta, 2\Delta, \ldots\}$ even if $\Delta > 0$ is arbitrarily small. Hence $\tilde{V}_t > \beta_t$ for all t, that is, S skims forever when the horizon is infinite and ρ is sufficiently close to zero. In particular this conclusion (skimming forever) applies for any $\Delta > 0$ and $\rho = 0$ (no discounting). This result has much in common with "no-trade theorems" familiar from the literature following Milgrom and Stokey (1982).

However, it is also true that for any given $\rho > 0$, there exists a sufficiently small period length ($\Delta > 0$ below a positive threshold) that induces S to *settle in finite physical time*. To see this note that

$$\lim_{\Delta \to 0} \tilde{V}_t = \frac{(r_B + r_S)\sigma}{r_B + r_S + \rho} \equiv \sigma\lambda < \sigma = \lim_{t \to 0} \beta(t).$$

Therefore, for any given $\rho > 0$, the seller won't skim forever provided $\Delta > 0$ is small enough, because $\tilde{V}_t > \beta_t$ fails for $t = \Delta > 0$ sufficiently close to zero.

Now fix $\rho > 0$ and assume $\Delta > 0$ is sufficiently small so that there indeed exists a final moment of settlement, $m > 0$. Both the time and the price of settlement depend on Δ; for a fixed $\rho > 0$, as $\Delta \to 0$, these quantities converge to

$$\tilde{m} = \frac{1}{r_S} \ln\left(\frac{\rho + (1-\sigma)(r_B + r_S)}{(1-\sigma)(r_B + r_S)}\right) \quad \text{and} \quad q(\tilde{m}) = \frac{(r_B + r_S)\sigma}{r_B + r_S + \rho} \equiv \sigma\lambda. \tag{14}$$

At any $t < m$ the seller may skim or settle at prices strictly less than 1 or β_t (respectively) because B anticipates settling at m. The surprising fact proven in the following Proposition is that for any given $\rho > 0$, if $\Delta > 0$ is arbitrarily small then S *settles right away* when the parties meet (in fact, even off the equilibrium path for all $t < m$).

Proposition 4 (Infinite Horizon: Equilibrium Settlement)**.** *Let* $\underline{t} = 0$ *and* $\bar{t} = \infty$*.*

For any fixed $\Delta > 0$*, if* $\rho \geq 0$ *is sufficiently small then in equilibrium* S *skims in every period with a price of 1, which* B *accepts iff either she has discovered or* S *has proved* $\omega = H$*.*

For any fixed $\rho > 0$*, if* $\Delta > 0$ *is sufficiently small then the equilibrium involves a last moment of settlement* $m > 0$*, moreover,* S *immediately settles (i.e., at* $t = \Delta$*) and does so even off the equilibrium path for all* $t \leq m$*, but skims forever for all* $t > m$*.*

For any fixed $\rho > 0$*, as* $\Delta \to 0$*, the last moment of settlement and the final settlement price converge to the quantities given in* (14)*. The settlement price for* $t \leq m$ *converges to* $q(t)$ *such that* $q(\tilde{m}) = \beta(\tilde{m}) = \sigma\lambda$*,* $q(t) < \beta(t)$ *for* $t < \tilde{m}$*, and for all* $t \in [0, \tilde{m}]$*,*

$$\dot{q}(t) = (\rho + r_B)[1 - \beta(t)]q(t) - (\rho + r_S)\beta(t)[1 - q(t)]. \tag{15}$$

If $\rho = 0$ *then, for any* $\Delta > 0$*,* S *skims forever with price* $p(t) \equiv 1$*.*

The case of skimming forever is a natural limit of the result described in the last section: Proposition 2 states that the last moment of settlement, m^*, is decreasing in the time horizon $\bar{t}$. As the horizon grows unboundedly, m^* shrinks towards zero, and eventually is smaller than Δ (if $\underline{t} = 0$). The trading parties have effectively "missed" the opportunity to reach agreement by the time they first meet, and the seller is sufficiently patient to wait for hard evidence to be discovered by either party and then sell at the maximal price.

The proof of the surprising bit (for any $\rho > 0$, settlement occurs at all $t = \Delta, \ldots, m$ provided Δ is sufficiently small) is given in the Appendix and is based on an induction argument: if $S_\emptyset$ anticipates that settling will be worthwhile a short time later at a price that $B_\emptyset$ will find acceptable (knowing that settlement will be offered in every period up until and including m), then it is optimal for $S_\emptyset$ to settle in the current period as well.

Here, the players' impatience is crucial: the basic idea's intuition given in Section 1 abstracts from impatience. The reason the seller wishes to delay before settling at m is the wedge between the seller's and buyer's pre-settlement beliefs. The importance of this wedge is eroded with impatience. As the length of the time period Δ shrinks, the chance either player receives a signal between times $m - \Delta$ and m shrinks with it. Eventually, for $\rho > 0$, impatience outweighs this chance, and the trading parties can find a mutually agreeable settlement price at $m - \Delta$ as well. Clearly this argument does not particularly depend upon m, and so by induction, the parties will wish to settle at all $t < m$.

There is no contradiction between the findings that

(i) for any fixed $\Delta > 0$, S strictly prefers to skim in all periods before m, provided $\rho \geq 0$ is sufficiently small (true no matter whether the time horizon is finite or infinite, see Proposition 1), and

(ii) for any fixed $\rho > 0$, S settles in all periods at and before m, provided $\Delta > 0$ is sufficiently small (true under infinite horizon, see Proposition 4).

First, the order of quantifiers is different in these two statements; second, with $\bar{t} \to \infty$ and $\rho = 0$ there is no final moment of settlement (S skims forever), hence there is only an apparent discontinuity of the latter result (settling at and before m) as ρ approaches 0.

The derivation of $\tilde{m}$ is illustrated in Figure 2, an infinite-horizon version of Figure 1. The uninformed seller's expected payoff from skimming forever is constant in time, converging to $\sigma\lambda$ as $\Delta \to 0$. The horizontal graph of $\tilde{V}$, in blue, intersects the graph of β, in red, exactly once, at $t = \tilde{m}$. The settlement price, $q(t)$ for $t \leq \tilde{m}$, is drawn in black. Settlement takes place immediately at the price given by $q(\underline{t})$.

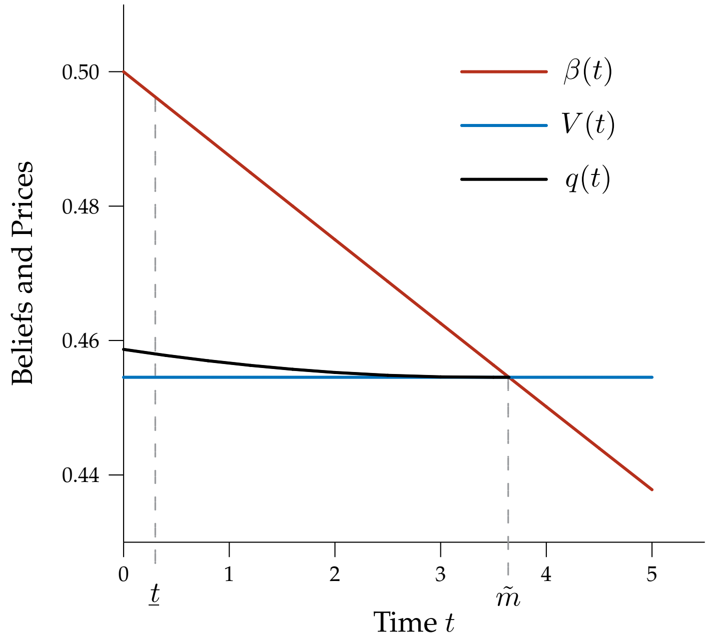

FIGURE 2. Immediate Settlement with an Infinite Horizon.

The figure illustrates a case of immediate settlement with an infinite horizon. The parameter values are $\sigma = 0.5$, $r_B = r_S = 0.05$, and $\rho = 0.01$. The uninformed seller's continuation value is constant in time; $\tilde{V} \to \sigma\lambda \approx 0.455$ as $\Delta \to 0$. $\beta(t)$ intersects $\tilde{V}$ exactly once at $\tilde{m} \approx 3.65$. Thus $q(\tilde{m}) = \beta(\tilde{m}) \approx 0.455$. Settlement takes place immediately at the price given by $q(\underline{t})$, and (off equilibrium path) a settlement price of $q(t)$ is offered at all subsequent $t \leq \tilde{m}$. After $\tilde{m}$, S skims with price 1.

This section is concluded with comparative statics results regarding the final settlement date and the price at which immediate settlement occurs with a fixed $\rho > 0$ and arbitrarily small $\Delta > 0$ in the infinite-horizon case. A greater common optimism about the good's value (a higher prior, σ) or more patience (lower discount rate ρ) increase the price at which the parties settle at the first trading opportunity.

Proposition 5 (Infinite horizon: Comparative Statics). *Let $\underline{t} = 0$, $\bar{t} = \infty$, fix $\rho > 0$, and let $\Delta > 0$ be arbitrarily small. The $\Delta \to 0$ limit of the final moment of settlement, $\tilde{m}$, is increasing in σ and ρ and decreasing in r_B and r_S. If $r_B = r_S = r$ then the price at which immediate settlement occurs, $q(\underline{t})$, is increasing in σ and decreasing in (a sufficiently low) ρ.*

The comparative statics of $\tilde{m}$ follow easily from inspecting (14). The proofs of the results on the immediate settlement price, $q(\underline{t})$, are more subtle and are relegated to the Appendix. For example, an increase in σ shifts both curves in Figure 2 up/out, and the new intersection occurs at a higher $\tilde{m}$ and $\beta(\tilde{m}) \equiv \sigma\lambda$; in the proof we then establish that all settlement prices before $\tilde{m}$ are higher as well.

A final observation is that since trade occurs immediately with a fixed $\rho > 0$ and arbitrarily frequent interaction, there is never any efficiency loss.

6. Concluding Remarks and Related Literature

The main, novel feature of the above model is that the parties may produce and selectively disclose verifiable evidence while participating in a dynamic trading game. As a result, the paper is related to the rich, but mostly distinct literatures on *communication games* (especially with "hard" information) and *dynamic bargaining models*.

Information structures with partially or fully verifiable evidence have mainly been applied in information transmission games rather than in bargaining (see Milgrom, 2008, for a survey). There, the central themes are whether communication results in the full revelation of all verifiable information (e.g., Seidmann and Winter, 1997; Mathis, 2008), and whether or not it is beneficial for a player to be known to be more able to generate hard evidence (e.g., Shin, 1994). One of the profound insights in this literature is that verifiability is a two-edged sword: it increases the credibility of the sender, but also the incredulity of the listener in case no message is forthcoming. This effect is clearly at work in the model here, albeit embedded in a model of dynamic price discrimination.

Persuasion as a technical term refers to communication via partially verifiable messages (e.g., Shin, 1994, mentioned above) and, relatedly, to the contractible design of a decision-maker's information structure (in the Bayesian persuasion literature, see Kamenica and Gentzkow, 2011). In our model both parties may be able to affect what the other knows (over time, by way of verifiable disclosure), but not without exogenous technological constraints, as evidence arrives stochastically and may only be either disclosed or concealed. Hence our game is less related to persuasion problems of the latter type.

There is a growing literature on dynamic information-transmission games, both with unverifiable messages (long cheap talk: Aumann and Hart, 2003), and using state-dependent message spaces (long persuasion: Forges and Koessler, 2008). In this line of research our motivation is perhaps closest to that of Hörner and Skrzypacz (2014) who study a specific persuasion technology (repeated, binary-outcome testing) in a dynamic game where communication and certain actions (payments) are concurrent. The sender-optimal equilibrium in their model exhibits a gradual release of information; in contrast, here (essentially by assumption) evidence disclosure immediately ends the game. The dynamic trading game in our paper is also very different from their sender-receiver game.

An early example of a trading model that involves hard information acquisition is that of Shavell (1994), where the focus is on the (in)efficiency of information acquisition rather than on the effects of verifiability itself. Our recent paper Eső and Wallace (2014) compares 'soft' and 'hard' evidence in a much simpler bargaining problem. In a single-period game with two-sided incomplete information, the equilibrium when the seller's information is unverifiable and the equilibrium when his information is verifiable are qualitatively different. Overall the seller is ex-ante better off when his signal is hard.

The dynamic trading game that we augment with the probabilistic acquisition and voluntary disclosure of hard evidence is familiar from the literature on dynamic monopoly and the Coase conjecture.[14] In these dynamic price-discrimination problems, as in ours, bargaining power is one-sided (i.e., the seller makes all the take-it-or-leave-it offers), and a central question is what happens as the period-length (representing the seller's ability to commit to a price) converges to zero. Recent papers by Fuchs and Skrzypacz (2010, 2013) study the stochastic arrival of new traders and deadlines in this type of model, whereas Daley and Green (2012, 2016) analyse the effects of the gradual, exogenous public revelation of one side's private information. In contrast to all of these analyses, in our model, both parties may acquire private information over time, and evidence is verifiable-concealable (it is neither "soft" private information nor public "news"). Despite the differences in modelling information structures, there are some shared conclusions, e.g., concerning the possibility of delay on the equilibrium path.[15] Comparative statics results are difficult to contrast because of the differences in modelling assumptions, but a common relationship emerges between prices and delay: longer delay (or a longer period of skimming before settlement) is associated with a lower settlement price.

Delay in bargaining has received much attention in earlier work. As is well known, delay may be caused by incomplete information (e.g., about the parties' valuations or discount factors), and either signalling (Admati and Perry, 1987) or a concern for reputation (Abreu and Gul, 2000). In the former, patient players signal their strength by refusing to agree for while; in the latter, a rational bargainer might strike a posture to appear as an "irrational type" by making demands and rejecting offers in order to build a reputation and get a better deal. Another intuitive reason for delay may be excess optimism (e.g., different priors about the distribution of bargaining power, or other behavioural anomalies), although Yildiz (2003) points out the limits of this type of explanation. Our model is related to this literature to the extent that it may also generate delay on the equilibrium path; the

[14]Seminal papers include Stokey (1981), Bulow (1982), Fudenberg, Levine, and Tirole (1985), Gul, Sonnenschein, and Wilson (1986), and Hart and Tirole (1988).

[15]Delay does not necessarily contradict the Coase conjecture (loss of excess profit in the no-commitment limit). For instance, in the model of Fuchs and Skrzypacz (2010), which exhibits equilibrium delay, the seller's ex-ante profit is the same as if he waited for his outside option to arrive. In contrast, Board and Pycia (2014) show the failure of the Coase conjecture in an environment where the party with an informational advantage has an outside option.

reasons (in our case, the concealment of hard evidence leading to the impossibility of mutually-acceptable trade) are however markedly different.

As alluded to earlier, the paper draws some of its inspiration from literary accounts of negotiations, where the emphasis is often on strategic evidence disclosure (persuasion) rather than the mechanics of offers, counteroffers, or deadlines. For example, Roald Dahl's gem of a story entitled *Parson's Pleasure* describes the adventures of a fictional but cunning London antiques dealer.[16] On Sundays he dresses up as a vicar (to disguise his expertise) and knocks on the doors of countryside homes under various pretenses; once he is in, he looks around for antiques to buy on the cheap. The heart of the story is a negotiation over a Chippendale commode. Our con-man/antiquarian tries to convince his uneducated hosts that the cupboard is a worthless fake, while they unwittingly produce irrefutable evidence of its provenance and exquisite value. In the end, he prevails: he pays twenty pounds for the item, and happily leaves to collect his car. The punchline of Dahl's story is that in the end, the fake vicar's persuasive arguments work too well. While he is away the hosts helpfully turn the commode into what he had convinced them it was worth: a pile of firewood. For a game theorist, what's remarkable is that most of the bargaining process concerns information transmission, whereas the price is set rather quickly. In fiction and in reality, the secret of success in negotiation is arguably the *production* and *disclosure* of credible evidence.

There are many different directions future research might take. One, mentioned earlier in Footnote 5, would be to consider imperfect hard evidence production in which there is some probability of receiving the high signal in the low state, and vice versa. A second would be to consider different trading environments (for instance, a different surplus in the two states or a continuum of states). Finally, all bargaining power is held by the seller in this model (to map more closely to the dynamic-trading literature discussed above). A different approach might model bargaining power shared between the trading parties; for example, via alternating offers à la Rubinstein (1982), or random offers.

Appendix A. Proofs and Derivations

Proof of Proposition 1. The existence and uniqueness of an equilibrium satisfying conditions (i) and (ii) were established (by construction) in the text. It remains to show that the seller skims for all $t < m$ provided $\rho > 0$ is sufficiently low. We prove this (with *strict* preference for skimming) for the case of $\rho = 0$; then the claim for ρ near zero follows by continuity.

[16]"Parson's pleasure" also refers to a secluded section on the River Cherwell where, in the distant past, Oxford dons would bathe in the nude. Our inspiration derives purely from Dahl's short story.

If at $t < m$ the next time of settlement (absent disclosure) is m at price $q_m = \beta_m$ then by (5) and (6), the skimming price is $p_t = 1 - (1 - r_S\Delta)^{(k-t)/\Delta}\beta_m$, and the settling price can be rewritten as

$$q_t = \left[1 - (1 - r_S\Delta)^{(m-t)/\Delta}\right]\beta_t + \left[(1 - r_S\Delta)^{(m-t)/\Delta}\beta_t + (1 - r_B\Delta)^{(m-t)/\Delta}(1 - \beta_t)\right]\beta_m. \tag{16}$$

Define V_t^m as the expected payoff of $S_\emptyset$ in period t conditional on no disclosure at or before t, from skimming in all periods from t leading up to m (including t but not m):

$$V_t^m = \left[1 - (1 - r_S\Delta)^{(m-t)/\Delta}\right]\hat{\sigma} + \left[(1 - r_S\Delta)^{(m-t)/\Delta}\hat{\sigma} + (1 - r_B\Delta)^{(m-t)/\Delta}(1 - \hat{\sigma})\right]\beta_m. \tag{17}$$

The seller gets a payoff of 1 if $\omega = H$ (i.e., with probability $\hat{\sigma}$) and he finds this out at or before m; whereas he gets β_m provided either the state is $\omega = H$ and he does not discover it or $\omega = L$ and B does not discover it by m.

By (16) and (17), skimming is more profitable for $S_\emptyset$ at t than settling, $V_t^m > q_t$, iff $\hat{\sigma} > \beta_t$. The latter, however, always holds by $\hat{\sigma} > \sigma = \beta_0 > \beta_t$ for all $t > 0$. □

Derivation of Formulæ for $\Delta \to 0$. Difference β_t given in (1) and divide by Δ to get

$$\frac{\beta_{t+\Delta} - \beta_t}{\Delta} = \frac{-r_S(1 - \beta_t)\beta_t}{(1 - r_S\Delta)\beta_t + (1 - \beta_t)}.$$

Take the $\Delta \to 0$ limit to get $\dot{\beta}(t) = -r_S\beta(t)[1 - \beta(t)] \in [-1, 0]$. Integrating this bounded derivative with boundary condition $\beta(0) = \sigma$ yields $\beta(t)$ as given in (7).

To derive $V(t)$ rearrange (3), cross-divide it by Δ, and take $\Delta \to 0$ to get

$$\dot{V}(t) \equiv \lim_{\Delta\to 0}\frac{V_{t+\Delta} - V_t}{\Delta} = (r_S + r_B + \rho)V(t) - (r_S + r_B)\sigma. \tag{18}$$

With boundary condition $V(\bar{t}\,) = \sigma v$ the closed-form solution in (8) is obtained.

The formula for m in the limit as $\Delta \to 0$, given in (4), follows from rearranging

$$\beta(m) = V(m) \equiv \left[1 - e^{-(r_B + r_S + \rho)(\bar{t} - m)}\right]\sigma\lambda + e^{-(r_B + r_S + \rho)(\bar{t} - m)}\sigma v.$$

To obtain $p(t)$ in (11) take the limit of p_t from (5) as $\Delta \to 0$ and set $k = m$.

Proof of Proposition 2. Note that changes in $\bar{t}$, v and ρ only affect V. It is immediate that an increase in $\bar{t}$ or v shifts V out (to the right or up), hence the locus of the last intersection between β and V (where both functions are decreasing, with β decreasing more rapidly) must decrease. Hence the final moment of settlement decreases and the corresponding settlement price increases (along β) for an increase in $\bar{t}$ or v. The skimming price in all periods before settlement decreases because $\frac{\partial}{\partial m^*}p(t) = [r_S + \rho - r_S\beta(m^*)]\,[1 - \beta(m^*)]\,e^{-(r_S+\rho)(m^*-t)} > 0$ by (11).

The opposite is true for an increase in ρ because (we claim) an increase in ρ reduces V for all $t < \bar{t}$, but not at the end of the time horizon. To see this suppose the opposite: that $V(t)$ remained the same or went up at $t < \bar{t}$ as ρ increased. Then by (18) the slope of V would become greater (less negative) at t, contradicting the requirement that V must remain fixed at $t = \bar{t}$. Hence an increase in ρ indeed rotates V counterclockwise (around its terminal value at $\bar{t}$), hence m^* increases, $\beta(m^*)$ falls, but $p(t)$ goes up for all $t < m^*$.

We claim $\partial m^*/\partial r_B < 0$. By inspecting $\dot{V}(t)$ in (18) and noting $V(t) < \sigma$ it is immediate that $\dot{V}(t)$ decreases (becomes more negative) for an increase in r_B, holding $V(t)$ fixed. Since $V(\bar{t}) \equiv \sigma v$ must remain the same even as r_B goes up, it follows that $V(t)$ must go up for all $t < \bar{t}$ for an increase in r_B. Since β does not depend on r_B it follows that the equilibrium level of m^* decreases. This implies $\beta(m^*)$ increases but $p(t)$ decreases for all $t < m^*$.

Next, $\partial m^*/\partial r_S < 0$ because an increase in r_S increases $V(t)$ for all $t < \bar{t}$ just like an increase in r_B does; moreover, a higher r_S also causes $\beta(t)$ to fall for all $t > 0$. Both effects result in a lower m^* and a higher $\beta(m^*)$.

Finally, to see $\partial m^*/\partial\sigma > 0$, first note that by $\beta(m^*)/\sigma = [\sigma + (1-\sigma)e^{r_S m^*}]^{-1}$ we have

$$\frac{d}{d\sigma}\left[\frac{\beta(m^*)}{\sigma}\right] = \left(\frac{\beta(m^*)}{\sigma}\right)^2 \left[e^{r_S m^*} - 1 - (1-\sigma)r_S e^{r_S m^*}\tfrac{\partial m^*}{\partial\sigma}\right].$$

Hence, by differentiating (9) in σ,

$$\begin{aligned}\frac{\partial m^*}{\partial\sigma} &= \frac{1}{r_B + r_S + \rho}\,\frac{\lambda - \beta(m^*)/\sigma}{\lambda - v}\left(\frac{\beta(m^*)}{\sigma}\right)^2 \left[e^{r_S m^*} - 1 - (1-\sigma)r_S e^{r_S m^*}\tfrac{\partial m^*}{\partial\sigma}\right] \qquad (19)\\ &= \frac{K(e^{r_S m^*} - 1)}{1 + K(1-\sigma)r_S e^{r_S m^*}} > 0,\end{aligned}$$

where K denotes the positive factor multiplying the bracketed expression in (19) for $\partial m^*/\partial\sigma$. The overall sign of $\partial m^*/\partial\sigma$ follows from $e^{r_S m^*} > 1$. □

Proof of Proposition 3. The results follow rather directly from (13).

Greater r_B or r_S increases λ, which is the relative weight on $1 > e^{-(r_B+r_S+\rho)m}$ in the bracketed term in (13), hence W increases holding m fixed. However, $\partial m/\partial r_i < 0$ for $i = B, S$, therefore $e^{-(r_B+r_S+\rho)m}$ goes up as well, increasing W even further.

An increase in ρ reduces λ, hence W as well, for a given m. However, m is increasing in ρ, further reducing W.

Both v and $\bar{t}$ reduce m (leading to an increase in W) but otherwise do not affect W directly.

Finally, the direct effect of σ on W is positive, but the indirect effect through increasing m is negative, hence the sign of $\partial W/\partial\sigma$ is ambiguous. □

Proof of Proposition 4. The claim that for any given $\Delta > 0$ and sufficiently small $\rho \geq 0$ the equilibrium prescribes for S to skim in every period was proved in the text.

It was also shown in the text that for any given $\rho > 0$ and sufficiently small $\Delta > 0$, in equilibrium, there is a last moment of settlement $m > 0$ where $\tilde{V}_m = \beta_m$. We now show by induction that for all $t < m$ the seller settles given any $\rho > 0$, for $\Delta > 0$ sufficiently small.

Write the uninformed seller's expected payoff at t, conditional on no disclosure at or before t, from skimming at t but then settling at price $q_{t+\Delta}$ in $t+\Delta$, as

$$V_t^{t+\Delta} = \hat{\sigma} r_B \Delta p_t + \hat{\sigma}\frac{(1-r_B\Delta)r_S\Delta}{1+\rho\Delta} + [\hat{\sigma}(1-r_B\Delta)(1-r_S\Delta)+(1-\hat{\sigma})(1-r_B\Delta)]\,\frac{q_{t+\Delta}}{1+\rho\Delta}.$$

The first term is S's payoff from B_H accepting $p_t = 1-(1-r_S\Delta)(1-q_{t+\Delta})/(1+\rho\Delta)$ as defined by (5); the second term is his payoff if S himself finds out $\omega = H$ within the current period and charges 1 next time. The third term is S's expected payoff from settling at t: he will get $q_{t+\Delta}$ if either $\omega = H$ but neither player finds it out, or $\omega = L$ and B does not learn it by next time.

After plugging in p_t and rearranging,

$$V_t^{t+\Delta} = \hat{\sigma}\left(r_B\Delta + \frac{1-r_B\Delta}{1+\rho\Delta}\right) - \frac{1-r_S\Delta}{1+\rho\Delta}\hat{\sigma}(1-q_{t+\Delta}) + \frac{1-r_B\Delta}{1+\rho\Delta}(1-\hat{\sigma})q_{t+\Delta}.$$

Next, the settlement price at t defined by (6) is

$$q_t = \beta_t - \frac{1-r_S\Delta}{1+\rho\Delta}\beta_t(1-q_{t+\Delta}) + \frac{1-r_B\Delta}{1+\rho\Delta}(1-\beta_t)q_{t+\Delta}.$$

Uninformed S wants to settle at t whenever $V_t^{t+\Delta} < q_t$, equivalently, after a cross-multiplication by $(1+\rho\Delta)$, cancelling terms and cross-dividing by $\Delta > 0$,

$$\hat{\sigma}\rho\Delta r_B + r_S\hat{\sigma}(1-q_{t+\Delta}) - r_B(1-\hat{\sigma})q_{t+\Delta} < \rho\beta_t + r_S\beta_t(1-q_{t+\Delta}) - r_B(1-\beta_t)q_{t+\Delta}.$$

After further rearrangement, $V_t^{t+\Delta} < q_t$ is equivalent to

$$(\hat{\sigma}-\beta_t)\left[r_B q_{t+\Delta} + r_S(1-q_{t+\Delta}) + \rho\right] < \rho\hat{\sigma}(1-r_B\Delta). \tag{20}$$

The proof is completed by observing that this condition indeed holds for any $t < m$ provided $\Delta > 0$ is sufficiently small.

Consider the limit of both sides of (20) as $\Delta \to 0$: if

$$[\sigma - \beta(t)]\left[r_B q(t) + r_S(1-q(t)) + \rho\right] < \rho\sigma \tag{21}$$

then (20) must also hold for $\Delta > 0$ sufficiently close to zero. The latter, however, follows because as $\Delta \to 0$, the final moment of settlement m converges to $\tilde{m}$ and price q_m converges to $\beta(\tilde{m}) = (r_B + r_S)\sigma/(r_B + r_S + \rho)$, hence $[\sigma - \beta(\tilde{m})](r_B + r_S + \rho) = \rho\sigma$. Then, by $q(t) \in (0,1)$, inequality

$$[\sigma - \beta(\tilde{m})]\left[r_B q(t) + r_S(1-q(t)) + \rho\right] < \rho\sigma$$

follows, hence the same is also true for m replacing $\tilde{m}$ in this inequality provided m is sufficiently close to $\tilde{m}$ (which it is, for $\Delta > 0$ sufficiently small). But then, for all $t < m$ we have $\beta(t) > \beta(m)$, hence (21) also holds. This completes the proof. □

Proof of Proposition 5. The fact that $\tilde{m}$ is increasing in σ and ρ but decreasing in r_B and r_S follow simply by inspecting (14), as mentioned in the main text.

In the rest of the proof assume $r_B = r_S \equiv r$. Recall that the settlement price at the final settlement date is $q(\tilde{m}) = \beta(\tilde{m}) = \sigma\lambda$, which is increasing in σ and decreasing in ρ because $\lambda = 2r/(2r+\rho)$. By (15), the slope of the settlement price is

$$\dot{q}(t) = (\rho + r)[q(t) - \beta(t)],$$

which is zero at $t = \tilde{m}$ and negative for all $t < \tilde{m}$, where $q(t) < \beta(t)$.

For clarity (in this paragraph only), denote the settlement price at t under prior σ by $q(t|\sigma)$ and the final settlement date by $\tilde{m}_\sigma$. Let $\sigma' > \sigma$. We know $\tilde{m}_{\sigma'} > \tilde{m}_\sigma$, moreover,

$$q(\tilde{m}_\sigma|\sigma') > q(\tilde{m}_{\sigma'}|\sigma') \equiv \sigma'\lambda > \sigma\lambda \equiv q(\tilde{m}_\sigma|\sigma),$$

where the first inequality follows by $\dot{q}(t|\sigma') < 0$ for $t < \tilde{m}_{\sigma'}$ and $\tilde{m}_\sigma < \tilde{m}_{\sigma'}$. Suppose towards contradiction that $q(t|\sigma') = q(t|\sigma)$ for some $t < \tilde{m}_\sigma$. But $\dot{q}(t|\sigma)$ is decreasing in the level of $\beta(t) \equiv \sigma/\left[\sigma + (1-\sigma)e^{rst}\right]$, which is itself increasing in σ; hence $\dot{q}(t|\sigma)$ is decreasing in σ, holding the level of $q(t|\sigma)$ fixed. Therefore $\dot{q}(t|\sigma') < \dot{q}(t|\sigma')$, which contradicts $q(\tilde{m}_\sigma|\sigma') > q(\tilde{m}_\sigma|\sigma)$. Therefore $q(t|\sigma') > q(t|\sigma)$ for all $t < \tilde{m}_{\sigma'}$, as claimed.

If $\rho = 0$ then the seller skims with $p(t) = 1$ for all $t > 0$; the value of the good is eventually revealed by either player so the expected present value of his payoff is the same as if he could settle at price σ at time 0 (which he cannot because $t = \Delta, 2\Delta, \ldots$). For $\rho > 0$ small the final settlement price is $\beta(\tilde{m}) < \sigma$, at $\tilde{m} \approx \rho/r$, which is obtained by Taylor expansion of (14) around $\rho = 0$. The immediate settlement price, $q(\underline{t})$ is less than σ as well, provided $\rho > 0$ is sufficiently small, since $\dot{q}(\tilde{m}) = 0$. □

References

ABREU, D., AND F. GUL (2000): "Bargaining and Reputation," *Econometrica*, 68(1), 85–117.

ADMATI, A. R., AND M. PERRY (1987): "Strategic Delay in Bargaining," *The Review of Economic Studies*, 54(3), 345–364.

AUMANN, R. J., AND S. HART (2003): "Long Cheap Talk," *Econometrica*, 71(6), 1619–1660.

AUSUBEL, L. M., P. CRAMTON, AND R. J. DENECKERE (2002): "Bargaining with Incomplete Information," in *Handbook of Game Theory with Economic Applications*, ed. by R. J. Aumann, and S. Hart, vol. 3, chap. 50, pp. 1897–1945. North-Holland, Amsterdam.

BOARD, S., AND M. PYCIA (2014): "Outside Options and the Failure of the Coase Conjecture," *American Economic Review*, 104(2), 656–671.

BULOW, J. I. (1982): "Durable-Goods Monopolists," *Journal of Political Economy*, 90(2), 314–332.

DAHL, R. (1979): "Parson's Pleasure," in *Tales of the Unexpected*, pp. 193–217. Penguin, London.

DALEY, B., AND B. GREEN (2012): "Waiting for News in the Market for Lemons," *Econometrica*, 80(4), 1433–1504.

——— (2016): "Bargaining and News," unpublished manuscript, Duke University and UC Berkeley.

ESŐ, P., AND C. WALLACE (2014): "Information and Evidence in Bargaining," *Economic Theory Bulletin*, 2(1), 23–32.

FORGES, F., AND F. KOESSLER (2008): "Long Persuasion Games," *Journal of Economic Theory*, 143(1), 1–35.

FUCHS, W., AND A. SKRZYPACZ (2010): "Bargaining with Arrival of New Traders," *American Economic Review*, 100(3), 802–836.

——— (2013): "Bargaining with Deadlines and Private Information," *American Economic Journal: Microeconomics*, 5(4), 219–243.

FUDENBERG, D., D. K. LEVINE, AND J. TIROLE (1985): "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information," in *Game-Theoretic Models of Bargaining*, ed. by A. E. Roth, chap. 5, pp. 73–98. Cambridge University Press, Cambridge.

FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, MA.

GUL, F., H. SONNENSCHEIN, AND R. WILSON (1986): "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 39(1), 155–190.

HART, O. D., AND J. TIROLE (1988): "Contract Renegotiation and Coasian Dynamics," *The Review of Economic Studies*, 55(4), 509–540.

HÖRNER, J., AND A. SKRZYPACZ (2014): "Selling Information," forthcoming, *Journal of Political Economy*.

KAMENICA, E., AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101(6), 2590–2615.

KENNAN, J., AND R. WILSON (1993): "Bargaining with Private Information," *Journal of Economic Literature*, 31(1), 45–104.

MATHIS, J. (2008): "Full Revelation of Information in Sender-Receiver Games of Persuasion," *Journal of Economic Theory*, 143(1), 571–584.

MILGROM, P. R. (2008): "What the Seller Won't Tell You: Persuasion and Disclosure in Markets," *Journal of Economic Perspectives*, 22(2), 115–131.

MILGROM, P. R., AND N. L. STOKEY (1982): "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26(1), 17–27.

MUTHOO, A. (1999): *Bargaining Theory with Applications*. Cambridge University Press, Cambridge.

RUBINSTEIN, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97–109.

SEIDMANN, D. J., AND E. WINTER (1997): "Strategic Information Transmission with Verifiable Messages," *Econometrica*, 65(1), 163–169.

SHAVELL, S. M. (1994): "Acquisition and Disclosure of Information Prior to Sale," *RAND Journal of Economics*, 25(1), 20–36.

SHIN, H. S. (1994): "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64(1), 253–264.

STOKEY, N. L. (1981): "Rational Expectations and Durable Goods Pricing," *The Bell Journal of Economics*, 12(1), 112–128.

YILDIZ, M. (2003): "Bargaining without a Common Prior—An Immediate Agreement Theorem," *Econometrica*, 71(3), 793–811.