Persuasion and Stubbornness in a Dynamic Trading Game*

Peter Eso† and Chris Wallace‡

Very preliminary, October 2011

Abstract

We propose a dynamic model of bilateral trade in which the parties can generate and verifiably disclose (or conceal, but not destroy) 'hard' signals about the good's value. We find that in equilibrium the seller may keep offering a high price, accepted only by a buyer who is concealing a good signal, until eventually settling on a lower price that is accepted for sure. During the delay an uninfomed buyer becomes increasingly convinced that the seller is concealing an unfavorable signal. The delay in trading (interpreted as stubbornness) is shorter if the seller is initially more optimistic (or the buyer more pessimistic) about the good's value, or if the time horizon is longer, or if the parties are more likely to receive a concealable hard signal.

*Preliminary version. We thank seminar participants at Oxford, University College London, Universidad Carlos III Madrid, Erasmus University Rotterdam, and in particular Thomas Noe for comments.

† Department of Economics, University of Oxford.

‡Department of Economics, University of Leicester. Corresponding author. Email: cw255@leicester.ac.uk.

1 Introduction

We propose a new model of bargaining where the parties can generate concealable hard signals over time. The hard signals, if obtained, can be used to persuade the other side and secure a more advantageous outcome. In this framework we investigate the price dynamics, the possibility of delay, and the overall efficiency of the equilibrium outcome. The model can accommodate commonly known but different priors about the amount of surplus to be split.[1] This makes it possible to address the role of optimism, pessimism and (rational) overconfidence in bargaining problems.

The existing and rather extensive literature on bargaining (modeled mostly as dynamic trading games) emphasizes the parties' asymmetries in 'soft' private information, impatience, and risk aversion.[2] We believe that in many real-world bargaining situations *persuasion* (the generation and disclosure of 'hard' signals) plays a similarly important role.[3] In financial disputes such as shareholder lawsuits the courts make their decisions (or settlements are reached) based on objective, hard evidence produced by the parties' experts. Hard evidence is generated by consultants (hired by all sides) in regulatory consultations as well.[4] Similarly, when pricing an initial public offering of shares (IPO), the seller and potential buyers do not merely engage in cheap talk regarding the firm's assets, or use costly signaling to reveal private information, but often try to generate hard evidence regarding expected future cash flows, market conditions, etc., in

[1] In contrast, Yildiz (2003) considers a model in which non-common priors are held over the probability of making future offers. Thanassoulis (2010) provides another model with non-common priors over the players' bargaining power.

[2] A seminal paper is Rubinstein (1981). Muthoo (1999) provides an overview of the literature.

[3] Optimal rules of persuasion are studied by Glazer and Rubinstein (2004).

[4] For example, in the UK, before the sale of telecommunication licences, Ofcom has sought input from potential market participants in the form of evidence-based analyses concerning auction rules and the effects of various other policy options.

order to persuade the other side and make the terms of the deal more attractive for themselves. Our model is intended to capture this feature of bargaining problems.

2 The model

We study a bilateral trading model with one seller (S, he), one buyer (B, she) and one, indivisible good. The buyer's valuation for the good is $v \in \{\underline{v}, \bar{v}\}$. The seller's prior belief about v is $\sigma_0 = \Pr_S(v = \bar{v})$, the buyer's is $\beta_0 = \Pr_B(v = \bar{v})$; the values of σ_0 and β_0 are commonly known.[5] We normalize $\underline{v} = 0$ and $\bar{v} = 1$ in order to simplify the formulas for expected values. We also assume that the seller's value for the good is 0, common knowledge. This is not a normalization: It expresses the assumption that trading is ex ante (weakly) Pareto efficient.[6]

The seller and buyer interact over periods $t = 1, \ldots, T$. The value of the good, v, does not change over time.

At the beginning of every period the seller may receive signal s_t, and the buyer signal b_t. With probability r_S the realization of s_t is v (i.e., the true value of the good, 0 or 1), otherwise $s_t = \emptyset$ (i.e., no signal). Similarly, $b_t = v$ with probability r_B, and $b_t = \emptyset$ otherwise. All signals are independent conditional on v across players and time. If $s_t = \emptyset$ then we may also say 'S did not observe v at t', likewise for B if $b_t = \emptyset$. We assume that a signal, if observed, can be verifiably disclosed. In other words, s_t and b_t are *hard signals*. A signal can also be concealed, but it cannot be destroyed. Moreover, a player cannot prove that he or she did not receive

[5] If $\sigma_0 = \beta_0$ then the players have common priors. If $\sigma_0 \neq \beta_0$ then they do not.

[6] A different assumption, leading to a different model, would be that the seller's value from not selling the good is c, where c may be positive or negative. If $c > 0$ then trading may not be efficient. If $c < 0$ then it is strictly efficient to trade in all circumstances.

a signal. The players decide simultaneously, right before the end of period t, whether or not to disclose any signal observed at or before t.

At the end of each period $t = 1, \ldots, T$ the seller proposes a price, p_t, and the buyer either accepts or rejects it. If she accepts then the parties trade; the seller's payoff is p_t and the buyer's is $v - p_t$, and the game ends. If no trade ever occurs then the payoffs are 0. Utilities are transferable, and both players are risk neutral. The seller discounts future (expected) payoffs using discount factor δ_S and the buyer does the same using δ_B; both δ_S and δ_B belong to $[0, 1]$.

To summarize, our dynamic trading game is played as follows. At every $t = 1, \ldots, T$ that is reached, verifiable signals $s_t = v$ and $b_t = v$ are generated independently and privately with probabilities r_S and r_B, respectively. The players simultaneously decide whether to reveal any signals that they have observed. Then, still in period t, the seller proposes p_t which the buyer either accepts or rejects. If p_t is accepted then the game ends with payoffs p_t for S and $v - p_t$ for B. Otherwise either the game continues with period $t + 1$ (if $t < T$) or ends with zero payoffs for both players (if $t = T$).

In the next section we derive the unique, pure-strategy perfect Bayesian equilibrium satisfying certain reasonable robustness criteria and discuss its properties.

3 Equilibrium

The solution concept that we use is perfect Bayesian equilibrium (PBE). In a game like ours PBE requires that (1) the players' strategies be best responses given their beliefs on and off the equilibrium path, and (2) the players' beliefs be consistent with their commonly-known priors and the equilibrium strategies in the usual sense. In our game, the players' priors regarding the probability of $v = 1$ are common knowledge, but may be different. As a consequence we need to keep track of the players' higher-order beliefs over the event $v = 1$.

In any equilibrium, if either player discloses a 1-signal then $v = 1$ becomes commonly known and the only price that S is willing to offer is 1. In the presence of discounting, by strict dominance, $s_t = 1$ must be immediately disclosed by the seller and trade takes place at $p_t = v$. We will simplify the exposition by assuming that $b_t = 0$ is also immediately disclosed and the parties trade at $p_t = 0$. (Whether a worthless good is traded at price 0 or not traded at all is immaterial.)

The equilibrium that we characterize has two additional properties that we mention in advance, mainly to better explain our notation. First, in our equilibrium an uninformed seller either sets $p_t = \bar{p}_t$ that is accepted for sure (i.e., the seller *settles*), or he sets $p_t = p'_t$, which is accepted if and only if the buyer has observed a 1 signal (i.e., the seller is *skimming*). Clearly, the seller could also set a price so high that no buyer would accept it — however we will show that this action is dominated by skimming.[7] As a result, an equilibrium price offer is rejected either because the buyer is uninformed or because she has made a mistake. The second property of the equilibrium concerns beliefs that are consistent with this observation: The seller believes on and off the equilibrium path that the buyer is uninformed if she has rejected his offer.[8]

It is important to note that these are properties of the equilibrium that we identify, not assumptions. From our constructive derivation it follows that there is a unique PBE with these properties.

[7]Yet another possibility is that a buyer made indifferent by a price offer accepts it at random. We focus on pure-strategy equilibria even though mixed strategies are allowed.

[8]For example, off the equilibrium path a seller cannot believe simultaneously that the buyer has observed a 1-signal and it is still possible for him to receive a 0. The qualitative features of our equilibrium would survive under alternative specifications of out-of-equilibrium beliefs.

We proceed by distinguishing two possible continuation equilibria after an offer is rejected in period t. In Case 1, the seller offers a skimming price in all periods $t+1,\ldots,T$. In Case 2, there is a subsequent period m where the seller settles at price p_m, whereas the seller sets a skimming price in all periods between t and m. In each case we shall derive the equilibrium price offer p_t and the buyer's equilibrium response. This will complete the derivation of the equilibrium.

Before starting the analysis it is useful to introduce two pieces of notation. Let σ_1 denote an uninformed seller's belief that $v=1$ conditional on the buyer being either uninformed or concealing a just-received 1-signal having not observed a signal until t:

$$\sigma_1 = \Pr_S\left(v=1 | b_t \in \{\emptyset, 1\}, b_\tau = \emptyset, \forall \tau < t\right) = \frac{\sigma_0}{\sigma_0 + (1-r_B)(1-\sigma_0)}.$$

This is the seller's belief at any t following the buyer's rejection of an offer. Let β_t denote the uninformed buyer's belief at t that $v=1$ conditional on the seller not having disclosed a 1-signal at or before t:

$$\beta_t = \Pr_B\left(v=1 | s_\tau \in \{\emptyset, 0\}, \forall \tau \leq t\right) = \frac{(1-r_S)^t \beta_0}{(1-\beta_0) + (1-r_S)^t \beta_0}.$$

Note that β_t is strictly decreasing in t. An uninformed buyer becomes increasingly suspicious that the seller who has not disclosed a 1-signal is concealing a 0-signal, making her less and less convinced that $v=1$.

3.1 Case 1: Skimming expected to the end

Suppose we are at $t=T$. By the construction of the equilibrium strategies and off-path beliefs, the seller believes the buyer is uninformed right before the beginning of the period. If the buyer remains uninformed at

T then she believes the good's value is $v = 1$ with probability β_t, which is therefore her maximum willingness to pay for the good. If no signal is revealed at T then an uninformed seller believes $v = 1$ with probability σ_1, and that the buyer has just generated a 1-signal at T, and hence is willing to pay 1 for the good, with probability $r_B\sigma_1$.

Therefore the seller offers either $p_T = 1$ (accepted only if $b_T = 1$, i.e., with probability $r_B\sigma_1$), or $p_T = \beta_T$ (accepted for sure). The seller sets $p_T = 1$ in the final period if and only if $r_B\sigma_1 \geq \beta_T$.

Let $V_T = r_B\sigma_1$, the seller's expected profit from skimming at T. If $V_T \geq \beta_T$, then as we have just argued the seller sets $p_T = 1$ and the buyer's continuation value at $T - 1$ is zero. Hence at $T - 1$, the seller can sell for sure at price $p_{T-1} = \beta_{T-1}$, or alternatively skim with $p_{T-1} = 1$. The seller could also *delay* (not sell for sure) by setting $p_{T-1} > 1$, but that can be shown to be inferior to skimming for any $\delta_S < 1$.

At any $t < T$: If for all $t + 1, \ldots, T$ the seller is expected to skim with a price of 1, then his payoff at t from skimming by setting $p_t = 1$ is

$$\begin{aligned} V_t = \sigma_1 r_B + \sigma_1(1 - r_B)\delta_S[1 - (1 - r_B)(1 - r_S)] + \ldots \\ + \sigma_1(1 - r_B)^{T-t}(1 - r_S)^{T-t-1}\delta_S^{T-t}[1 - (1 - r_B)(1 - r_S)]\}. \end{aligned}$$

The first term of this expression is the probability that the buyer is concealing a 1-signal times the skimming price (which is 1). The second term is the probability that the buyer is not concealing a 1-signal but in the next period either the seller or the buyer will generate a 1-signal, which is again multiplied by the trading price (of 1) and discounted at rate δ_S. The remaining terms are calculated similarly.

This finite geometric series can be summed as

$$V_t = \sigma_1 \left\{ r_B + [1-(1-r_B)(1-r_S)] \times \frac{(1-r_B)\delta_S \left[1-(1-r_B)^{T-t}(1-r_S)^{T-t}\delta_S^{T-t}\right]}{1-(1-r_B)(1-r_S)\delta_S} \right\}$$

For $\delta_S = 1$, V_t simplifies to

$$V_t = \left[1-(1-r_B)^{(T-t)+1}(1-r_S)^{T-t}\right]\sigma_1.$$

Note that this is just the probability that either the buyer is already concealing a 1-signal or will get such a signal in the future or the seller will get a 1-signal in the future. If this event occurs the seller will sell the good at price 1, otherwise he will not be able to sell (as he is skimming by setting a price of 1 in all periods). Due to the lack of discounting the seller's profit is the expected price at which he (eventually) sells.

For any $\delta_S \in [0,1]$, it is straightforward to verify that V_t is decreasing in t, increasing in T, r_S, r_B, and σ_1; at $\delta_S = 1$ we have $\lim_{T\to\infty} V_t = \sigma_1$. We plot V_t for $\delta_S = 1$ in Figure 1. The following Lemma establishes exactly when Case 1 (skimming in all periods) applies, either from the beginning of the game or in any subgame. Lemma 1 holds for any $\delta_S \in [0,1]$.

Lemma 1: If $V_t \geq \beta_t$ for all t, then there is an equilibrium in which

(i) $s_t = 1$ or $b_t = 0$ is disclosed, followed by trade at $p_t = v$;

(ii) for all t without such disclosure the seller sets $p_t = 1$, and the buyer buys the good as soon as she observes a 1-signal.

3.2 Case 2: Settling expected at $p_m = \bar{p}$

If $V_T \equiv r_B\sigma_1 < \beta_T$ then the seller settles in period T, at price $p_T = \beta_T$. (That is, unless a signal is disclosed at T, the parties trade at price β_T.) In what follows we characterize the seller's optimal decision at t if he is expected to settle in period $m > t$, at price $p_m = \bar{p}$, but skim in all periods between t and m.

If $m > t+1$ then the seller is skimming in period $m-1$. A buyer who has seen a 1-signal knows that she can reject the skimming price at $m-1$ and buy the good a period later at price p_m, unless the seller also discovers a 1-signal by then. Therefore the informed buyer's continuation value at $m-1$ is $\delta_B(1-r_S)(1-\bar{p})$. The highest price that she accepts satisfies $1 - p'_{m-1} = \delta_B(1-r_S)(1-\bar{p})$, which pins down p'_{m-1}. By the same argument, for all $i = 0, 1, \ldots, m-t-1$, the skimming price at $t+i$ is

$$p'_{t+i} = 1 - \delta_B^{m-t-i}(1-r_S)^{m-t-i}(1-\bar{p}).$$

If an uninformed buyer rejects the seller's offer at t then her continuation value is

$$\begin{aligned} U_t = {} & \delta_B\beta_t(1-r_S)r_B(1-p'_{t+1}) + \delta_B^2\beta_t(1-r_S)^2(1-r_B)r_B(1-p'_{t+2}) \\ & \ldots + \delta_B^{m-t-1}\beta_t(1-r_S)^{m-t-1}(1-r_B)^{m-t-2}r_B(1-p'_{m-1}) \\ & + \delta_B^{m-t}\left[\beta_t(1-r_S)^{m-t}(1-r_B)^{m-t-1}(1-\bar{p}) - (1-\beta_t)(1-r_B)^{m-t}\bar{p}\right]. \end{aligned}$$

Here the first term is the discounted payoff the buyer receives when, in the next period, she is skimmed at price p'_{t+1}, which happens only when the buyer observes a 1-signal but the seller does not. All except the last term are analogous. The final term is the discounted expected payoff from the settlement price, $\bar{p}$ in period m.

After substituting in the formula for p'_{t+i}, for $i = 1, \ldots, m-t-1$, this simplifies to

$$U_t = \delta_B^{m-t}\left[(1-r_S)^{m-t}\beta_t(1-\bar{p}) - (1-r_B)^{m-t}(1-\beta_t)\bar{p}\right].$$

The only positive surplus that an uninformed buyer can receive obtains in the event that the seller remains uninformed when $v = 1$; in this case she pays only $\bar{p}$ in period m but receives a good of value 1. However, an uninformed agent obtains a negative surplus in the event that $v = 0$ and she remains uninformed about it until m; in this case she pays $\bar{p}$ at m for a worthless good. An uninformed buyer in period t accepts any price $p_t \leq \beta_t - U_t$.

It is easy to show that the maximum price accepted by a buyer concealing a 1-signal, $p'_t = 1 - \delta_B^{m-t}(1-r_S)^{m-t}(1-\bar{p})$, exceeds that accepted by an uninformed buyer, $\bar{p}_t = \beta_t - U_t$. Therefore, at t the seller has three choices: (1) Offer p'_t to *skim* the buyer concealing a 1-signal; (2) Offer $\bar{p}_t$, accepted for sure, to *settle*; (3) Offer $p_t > p'_t$ rejected for sure, to *delay*. It can be shown that *skim* dominates *delay* for all $\delta_S < 1$. This is so because delay forgoes profit on the buyer currently concealing a 1-signal.

The seller's payoff from skimming at t is

$$\begin{aligned} V_t^m = {} & \sigma_1 r_B p'_t + \sigma_1(1-r_B)\delta_S[r_S + (1-r_S)r_B p'_{t+1}] + \dots \\ & + \sigma_1(1-r_B)^{m-t-1}(1-r_S)^{m-t-2}\delta_S^{m-t-1}[r_S + (1-r_S)r_B p'_{m-1}] \\ & + \sigma_1(1-r_B)^{m-t}(1-r_S)^{m-t-1}\delta_S^{m-t}[r_S + (1-r_S)\bar{p}] \\ & + (1-\sigma_1)(1-r_B)^{m-t}\delta_S^{m-t}\bar{p}. \end{aligned}$$

The first term of this sum is the expected profit from skimming with price p'_t at t; this price is accepted only when the buyer is concealing a 1-signal. If that offer is rejected then the seller makes a sale in the following period if either he is able to reveal a 1-signal (in which case the price is 1), or he does not observe a 1-signal but the buyer does and accepts his skimming price at $t+1$. The remaining terms are analogous; the final two terms represent the expected profit from selling in period m, when the seller settles at $\bar{p}$ unless he receives a 1-signal. This series can also be written in a more compact form. In particular, for $\delta_S = \delta_B = 1$ the formula simplifies to

$$V_t^m = \left[1 - (1-r_S)^{m-t}\right]\sigma_1 + \left[(1-r_S)^{m-t}\sigma_1 + (1-r_B)^{m-t}(1-\sigma_1)\right]\bar{p}.$$

Skimming is better than settling for the seller at t if and only if $V_t^m \geq \bar{p}_t$. The comparison is particularly revealing when $\delta_S = \delta_S = 1$. For $\delta_B = 1$ the settling price $\bar{p}_t = \beta_t - U_t$ becomes

$$\begin{aligned}\bar{p}_t &= \beta_t - \left[(1-r_S)^{m-t}\beta_t(1-\bar{p}) - (1-r_B)^{m-t}(1-\beta_t)\bar{p}\right] \\ &= \left[1-(1-r_S)^{m-t}\right]\beta_t + \left[(1-r_S)^{m-t}\beta_t + (1-r_B)^{m-t}(1-\beta_t)\right]\bar{p}.\end{aligned}$$

Therefore, if $\delta_S = \delta_S = 1$, then $V_t^m \geq \bar{p}$ is equivalent to $\sigma_1 \geq \beta_t$. Note that the condition does not depend on the value of the settlement price, $\bar{p}$, nor the period when the settlement is expected to occur.

We summarize our findings in Case 2 (settlement expected in period m, at price $\bar{p}$) as follows.

Lemma 2: Let m be the greatest $t \leq T$ with $\beta_t > V_t$.

(i) Absent signal disclosure at or before m, trade occurs at $p_m = \beta_m$.

(ii) Absent signal disclosure at or before $t < m$, the uninformed seller skims at t with price $p_t' = 1 - (1-r_S)^{m-t}(1-\beta_m)^{m-t}$, if and only if $V_t^m \geq \beta_t - U_t$. Otherwise the seller settles at t, price $\bar{p}_t = \beta_t - U_t = \beta_t - \delta_B^{m-t}\left[(1-r_S)^{m-t}\beta_t(1-\beta_m) - (1-r_B)^{m-t}(1-\beta_t)\beta_m\right]$.

(iii) For $\delta_S = \delta_B = 1$, the condition $V_t^m \geq \beta_t - U_t$ is equivalent to $\sigma_1 \geq \beta_t$. Hence for discount factors close to one, the seller skims at $t < m$ whenever $\sigma_1 > \beta_t$.

3.3 Characterization of the equilibrium

The analysis so far can be combined in the following Theorem. A key, qualitative feature of the equilibrium is that, absent signal disclosure, the seller offers a price (not necessarily 1) only acceptable to a buyer concealing a 1-signal for some period of time. After this period of time he offers a price that is accepted for sure. If that is rejected (off the equilibrium path) then the seller reverts to skimming in all subsequent periods.

We first state the Theorem and then illustrate its contents in the discussion that follows.

Theorem 1. For δ_S, δ_B close to 1, the equilibrium is as follows:

1) If $\beta_t < V_t$ for all $t = \tau, \ldots, T$ then the seller offers $p_t = 1$ for all $t \geq \tau$. Trade occurs at a positive price if and only if a 1-signal is generated at t.

2) Suppose there exists $m = \max\{t \leq T : \beta_t \geq V_t\}$.

Assume $\sigma_1 > \beta_1$. There is no trade before m unless either the seller observes $s_t = 1$ (disclosed, trade at $p_t = 1$), or the buyer observes $b_t = 0$ (also disclosed, trade at $p_t = 0$), or the buyer observes $b_t = 1$ (concealed) in which case she accepts $p_t = 1 - \delta_B^{m-t}(1 - r_S)^{m-t}(1 - \beta_m)$. In all other cases trade takes place at m, price $p_m = \beta_m$.

If $\sigma_1 < \beta_1$ then, in the absence of signal disclosure at $t = 1$, trade occurs at $p_1 = \beta_1 - \delta_B^{m-1}\left[(1 - r_S)^{m-1}\beta_1(1 - \beta_m) - (1 - r_B)^{m-1}(1 - \beta_1)\beta_m\right]$.

As we discussed at the beginning of the section, if either S generates $s_t = 1$ or B generates $b_t = 0$, then the signal and hence the value of the good is revealed and trade occurs immediately at $p_t = v$. The main question that Theorem 1 answers is whether in a given period trade also occurs *without* signal disclosure (when either the players observed no signals, or one of them observed a disadvantageous signal and concealed it). The main result of the theorem is that for discount factors close to 1, there are essentially three possibilities, depending on the parameter values.

---

1) If $\beta_t < V_t$, for all t, then skim with $p_t = 1$ to the end:

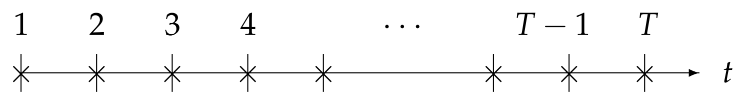

2A) Otherwise, if $\beta_1 < \sigma_1$ then skim for all $t < m = \max\{t | \beta_t > V_t\}$ and settle at $p_m = \beta_m$:

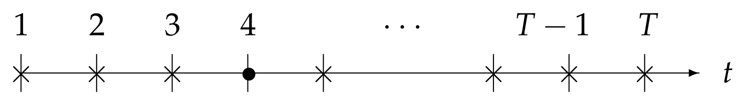

2B) If $\beta_t > V_t$ for some t and $\beta_1 > \sigma_1$, then settle at $t = 1$ and at all $t < m$ such that $\beta_t > \sigma_1$, then settle again at m, skim otherwise:

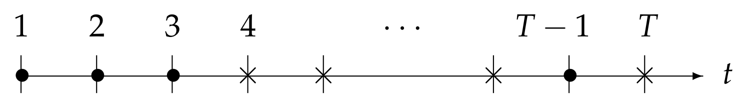

We illustrate Case 2A graphically in Figure 1. In this example,

$$\sigma_1 > \beta_1 > V_1 \equiv [1 - (1 - r_B)^T (1 - r_S)^{T-1}]\sigma_1 > V_T \equiv r_B \sigma_1 > \beta_T.$$

The first inequality, $\sigma_1 > \beta_1$, means that the seller is more enthusiastic about the good's value than the buyer at the time they meet. (This would be the case with common priors but it can also happen otherwise.) Then, provided $\beta_1 > V_1 \equiv [1 - (1 - r_B)^T (1 - r_S)^{T-1}]\sigma_1$, trade in the absence of signal realization takes place at $m = \max\{t \leq T : \beta_t > V_t\}$. Assuming $V_T \equiv r_B \sigma_1 > \beta_T$ such $m < T$ exists.

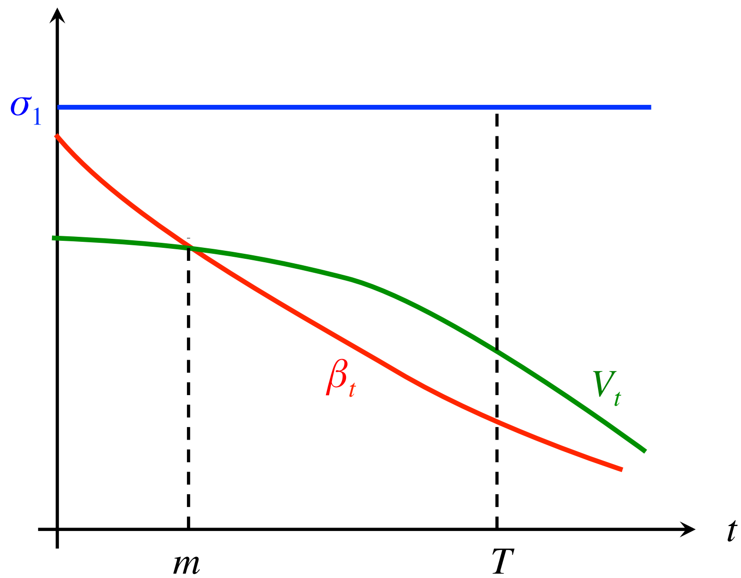

Figure 1: Illustration of an equilibrium with trade at $m \in (1, T)$.

4 Discussion and applications

In this section we explore the consequences of the main characterization theorem in a variety of settings. First, we perform some comparative static excercises on the length of time until trade will occur with probability one. Then we examine three special cases that highlight certain implications of the model.

4.1 Comparative statics of m in the parameters

Recall that m is the last period in which $\beta_t \geq V_t$ hence from the Theorem, if m exists then it is the last period in which trade occurs with certainty (i.e., even in the absence of signal disclosure). The following Proposition describes how this "last period of compromise" changes as a function of the parameters of the model.

Proposition 1: (i) If the seller is a priori more optimistic, or if the buyer is less optimistic, then m is lower (compromise is reached earlier).

(ii) A larger T (longer horizon) shifts V_t right inducing a lower m.

(iii) Larger r_S or r_B shifts V_t up, β_t down (constant in r_B), decreasing m.

Notice that this Proposition has particular impact when Case 2A obtains ($m-1$ periods of skimming followed by settling at $p_m = \beta_m$). Of course, it also affects the equilibrium price in Case 2B where the parties settle immediately.

4.2 Discontinuity of the equilibrium mapping

Assume δ_B and δ_S are both close to 1. Start from a situation where $\sigma_1 > \beta_1$ (e.g., common prior), and $\beta_T > V_T$ (e.g., r_B not too large). Then the equilibrium outcome is for the seller to skim for all $t < T$ and settle at $p_T = \beta_T$.

Suppose the seller becomes somewhat skeptical relative to the buyer, i.e., σ_0 is decreased while β_0 remains the same. All else equal, σ_1 as well as V_t for all t decrease.

If, as a result of this change in the prior, σ_1 falls just below β_1, then the outcome changes drastically. The parties immediately settle at price $p_1 = \beta_1 - \delta_B^{T-1}\left[(1-r_S)^{T-1}\beta_1(1-\beta_T) - (1-r_B)^{T-1}(1-\beta_1)\beta_T\right]$.

In this example, as σ_0 decreases, the equilibrium strategies (and the corresponding outcome) may change discontinuously from *'skim throughout'* to *'settle right away'*.

4.3 The potential benefits of excessive seller enthusiasm

Consider an example with a common prior, $T = 1$, but assume the seller can make an *ex ante offer* p_0. Assume $\beta_1 > V_1$.

Since $\sigma_0 = \beta_0$ we have immediate agreement at time 0. Off the equilibrium path the parties settle at $p_1 = \beta_1$. Hence $p_0 = r_S\sigma_0 + (1 - r_S)\beta_1$.

Suppose that the seller hires an agent with $\sigma_0' > \sigma_0$ such that $V_1' = r_B\sigma_0' > \beta_1$. At $t = 1$ the seller would skim with $p_1' = 1$. Anticipating this, the settlement price at 0 is $\beta_0 = \sigma_0 > p_0$. In this example the seller can gain from pretending to be overly enthusiastic because this commits him to delay off the equilibrium path.

4.4 Large, finite T and discounting

Suppose $\sigma_1 > \beta_1$ (e.g., common prior), and let $\delta_B = \delta_S = \delta \in [0, 1]$, with T arbitrarily large. Then $V_t \approx V(\delta)$ which is constant for all t. Note $V(0) = r_B\sigma_1$, $V(1) = \sigma_1$, and $V(\delta)$ is strictly increasing in δ. In contrast, $\beta_T \to 0$ as $T \to \infty$. Therefore we have the following result.

Proposition 2. Assume $\sigma_1 > \beta_1$, $\delta_B = \delta_S = \delta \in [0, 1]$, and T large.

(i) If $\delta \approx 0$ then in equilibrium trade occurs without signal generation at the greatest integer $m(0) < T$ such that $\beta_{m(0)} \geq r_B\sigma_1$.

(ii) As δ increases, trade without signal generation occurs at $m(\delta)$, which is decreasing in δ. For δ near 1, there is no trade without signal generation.

The proposition is illustrated in Figure 2. Interestingly, for low values of δ, as the players become more patient (δ goes up), they reach a compromise earlier. This may seem counterintuitive. The explanation is that as δ rises, the value to the seller from skimming in all future periods (V_t) increases. The seller can offer a compromise price when the uninformed buyer's willingness to pay (here, β_t) first dips below the seller's continuation value. Since β_t is decreasing over time, this happens sooner when V_t is uniformly greater for all t. However, for δ sufficiently close to 1, compromise is never reached (i.e., the seller skims in all periods). This latter fact is straightforwardly explained: because the seller is initially more confident than the buyer ($\sigma_1 > \beta_1$), he is prepared to wait indefinitely for a 1-signal to be revealed.

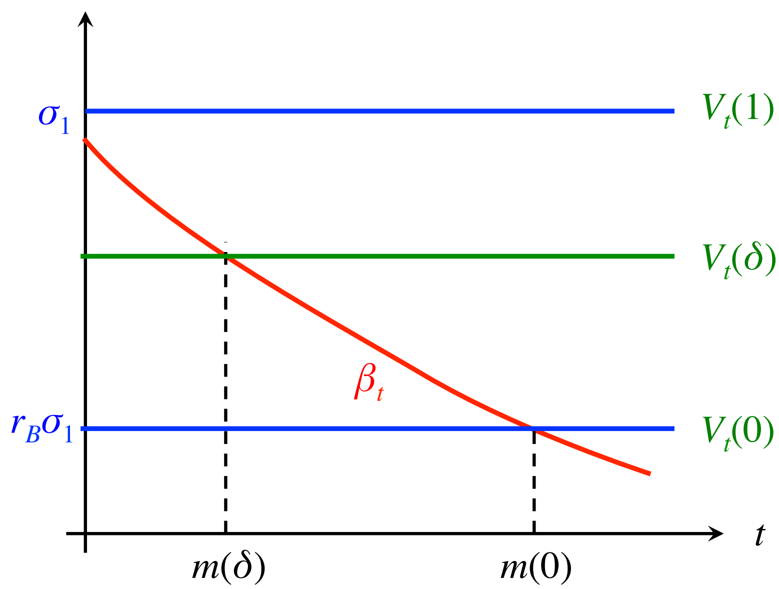

Figure 2: Long horizon and discounting.

5 Related literature

[to be added]

6 Conclusions

We have developed a model of bilateral trading with persuasion (generation of concealable hard signals) and possibly different, commonly-known priors. The equilibrium exhibits stubbornness: the seller offering a high price rejected by an uninformed buyer who grows more and more suspicious over time. However, compromise may be reached when the seller becomes sufficiently worried that no sale will take place at all.

We examined the effects of discounting, the probability of generating hard signals, and the length of the time horizon on the date of compromise. We also looked at the way in which the buyer's and seller's priors (skepticism or optimism) affect the equilibrium outcome.

There is a number of interesting avenues for further research, including the extension of the model to allow ex-ante trade, with or without a common prior; the incorporation of a more general offer structure to represent bargaining power (e.g., alternating or random offers). A similar structure with concealable hard information could also be incrorporated in other, related models of trade and bargaining, for example, ones with the possibility of a negative surplus, auctions, and debate.

References

Glazer, Jacob and Ariel Rubinstein (2004), "On optimal rules of persuasion", *Econometrica*, 72(6): 1715-1736.

Muthoo, Abhinay (1999), *Bargaining Theory with Applications,* CUP, Cambridge.

Rubinstein, Ariel (1982), "Perfect equilibrium in a bargaining model," *Econometrica*, 50(1): 97-109.

Thanassoulis, John (2010), "Optimal stalling when bargaining," *Journal of Economic Dynamics and Control*, 34:101-120.

Yildiz, Muhamet (2003), "Bargaining without a common prior—an immediate agreement theorem," *Econometrica*, 71(3): 793-811.