RESEARCH PAPER NO. 1765

Uncertainty about Uncertainty and Delay in Bargaining

Yossi Feinberg
Andrzej Skrzypacz

September 2002

RESEARCH PAPER SERIES

STANFORD
GRADUATE SCHOOL OF BUSINESS

Uncertainty about Uncertainty and Delay in Bargaining

Yossi Feinberg and Andrzej Skrzypacz
Stanford University, Graduate School of Business

September 19, 2002

Abstract

We study a one-sided offers bargaining game in which the buyer has private information about the value of the object and the seller has private information about his *beliefs* about the buyer. We show that this uncertainty about uncertainties dramatically changes the set of possible outcomes when compared to two-sided private information. In particular, higher order beliefs can lead to a delay in reaching agreement even when the seller makes frequent offers, while in the case of two-sided first order private information, agreement is reached almost instantly. Furthermore, we show that not all types of higher order beliefs lead to a delay: the crucial condition is that when uncertain about uncertainties, one assigns positive probability to *certainty*.

1 Introduction

The inclusion of uncertainty and private information in theories of bargaining, auctions, entry, principle-agent and many other topics, has been instrumental in our analysis and understanding of these strategic interactions. However, these models are usually of the following form: Some of the participating players may have private information about some fundamental uncertainty (usually relating to the players' payoff) and any uncertainty about these fundamentals by the other players is assumed to be commonly known. Consider a bargaining situation where the seller has an object that the potential buyer values. The fundamentals are the valuation that the buyer assigns to the object and the costs to the seller for providing the object (if any). Throughout the vast literature about bargaining, the exact uncertainty that the seller or the buyer might have about these fundamentals is itself assumed to be commonly known. For example, the seller may be uncertain as to the buyer's valuation but the beliefs that she holds about these valuations are assumed to be commonly known.

In this paper we relax this restriction. We postulate that if a seller is uncertain as to a buyer's valuation, it seems quite restrictive to assume that the buyer can figure out the exact beliefs that the seller possess. Why should the buyer be able to read the seller's mind as to her subjective beliefs when the seller cannot read the buyer's mind as to his valuation? While one can come up with a justification for knowledge of others' uncertainties in some cases[1], we claim that relaxing this constraint merits consideration.

The question is whether the consideration of higher order uncertainties provides new insights beyond those provided by the prevailing models. In this paper we demonstrate that uncertainties about uncertainties have a major impact on equilibrium behavior in a bargaining framework. We analyze the extent of this impact and its origins. More specifically, we show that uncertainties about uncertainties may lead to a delay in reaching an agreement in a bargaining situation with frequent one-sided offers, even when there is common knowledge of positive gains from trade.

This result is in contrast to the well-established Coase property for models with one-sided uncertainty about fundamentals and one-sided frequent offers, where common knowledge of gains from trade implies immediate agreement. To isolate the impact of higher order beliefs on delay, we also show that these results are not a consequence of *two*-sided uncertainty. We achieve delay within a class of sequential equilibria that include many of the equilibria constructed in the reputation literature (cf. Kreps and Wilson 1982, Milgrom and Roberts 1982 and Abreu and Gul 2000). We show that common knowledge a *large* gap between the seller's costs and the buyer's valuation implies that there is no delay in the standard two-sided private information. However, the main contribution is that with the same common knowledge of a large gap we do get delay in agreement when *high order* uncertainties

[1] One could think of a seller and a buyer being chosen from a large population. If the distribution of valuations and costs in the population is of public record, then this objective distribution generates the commonly used two-sided private information framework. However, assuming any uncertainty or private information as to the distribution of the population translates into the situation discussed in this paper.

are present[2]. Considering the sometimes elaborate attempts necessary for achieving delay in bargaining, these results cast a shadow on the robustness of the prevailing assumption of common knowledge of first order beliefs.

Consider a buyer and a seller bargaining. Assume that the seller is making offers which the buyer can accept or reject. The seller's cost is commonly known – assume that her cost is zero. The buyer's valuation may only be known to himself. As long as the uncertainty about the buyer's private information is commonly known, in the "gap case" (the buyer's lowest possible valuation is strictly positive) agreement is achieved instantly as offers are made more and more frequently. This well established result by Fudenberg, Levine and Tirole (1985) and Gul, Sonnenschein and Wilson (1986) has motivated the consideration of variations of bargaining procedures which may lead to delay in bargaining – a delay that is so readily observed in actual bargaining. We show that merely the buyer's uncertainty regarding the seller's beliefs about his valuation is enough for a whole new set of equilibria to emerge. These equilibria all lead, with positive probability, to a delay in agreement even when the seller makes frequent offers.

It is important to note that the thrust of our result is not the mere demonstration of the failure of the Coase property, but rather the impact of higher order uncertainty. Delay in bargaining was shown to occur under two-sided alternating offers in Admati and Perry (1987), when no gap exists in a two-sided private information scenario in Cramton (1992) (see also Cramton 1984), when cost and value are correlated and the gap is small in Vincent (1989) or when irrational types are present in Abreu and Gul (2000). It is the contrast between uncertainty about fundamentals and uncertainty about beliefs that drives our results.

It turns out that not every uncertainty about beliefs leads to delay in bargaining in the scenario mentioned above. For example, if the buyer could have either a low or a high strictly positive valuation and the seller has a cost of zero, then delay occurs *only* when the buyer deems it possible that the seller might be *certain* of his (the buyer's) valuation. Put another way, the buyer may be uncertain as to the seller's beliefs, but he must assign positive probability that she actually knows what his valuation is for delay to occur. However, if the seller's possible types always assign positive probabilities to

[2]We are focusing on a *large* gap since Evans (1989), Vincent (1989), and Deneckere and Liang (1999) have shown that in models with correlated private information on both sides, common knowledge of gains from trade is not sufficient to guarantee no delay in equilibrium. However, when the gap is large enough, then the no-delay result holds.

each and every possible valuation of the buyer (full support) then no delay will occur. We call this condition the *possible exclusion of types*: it must be possible for the seller to have a type that excludes some of the buyer's types for delay to emerge. We further discuss this property in the last section of the paper.

Our analysis follows the tools laid by Kreps and Wilson (1982) for games with two-sided private information. The principle is closely related to the reputation literature. In a nutshell, both the buyer and the seller have a type they would like to mimic when they are of the other type: the seller would like to mimic the type that is informed of the buyer's high valuation and the buyer would like to mimic the low valuation type. Such a scenario allows for the new type of sequential equilibrium behavior to emerge. It also explains the necessity of the condition of possible exclusion of types, since if the condition fails, *all* of the seller's types have full support over the buyer's types. This implies that even if the buyer knew the seller's type, he would have expected a unique Perfect Bayesian equilibrium behavior (in the limit) – that is, identical behavior by all of the seller's types, so there is no incentive to mimic a specific type.

The equilibria we construct are stationary: the strategies depend only on the current state of the game which is defined by the current beliefs. It makes the equilibrium delay even more striking, since in the existing literature stationarity is found to drastically improve efficiency of trade. For example, even in the no-gap case with one-sided private information stationarity implies no delay (see Fudenberg, Levine and Tirole 1995 and Gul, Sonnenschein and Wilson 1986). Furthermore, Cho (1990) shows that stationarity guarantees no delay in a model with two-sided private information about fundamentals with strictly positive gains from trade.

In section 2 we present the model. Section 3 contains the main result where uncertainty about uncertainties is incorporated into the model leading to delay in bargaining. In section 4 we discuss and prove the necessity of the possible exclusion of types condition and discuss further variations. The last section 5 discusses possible extensions and additional remarks.

2 The model

We consider a buyer and a seller who bargain over a sale of one item. It is common knowledge that the seller values retaining the object at zero, or, equivalently that the cost to the seller for producing the object is zero. The value to the buyer is either h or l with $h > l > 0$. We will consider various information structures relating to uncertainties about the fundamentals h and l, but we will always assume that the buyer knows his value and that this fact is commonly known. The seller could be uncertain as to the buyer's valuation. Such an uninformed seller will have initial beliefs characterized by the probability α that he assigns to the buyer having the valuation h with $0 \leq \alpha \leq 1$. These initial beliefs will characterize the possible type of seller. Finally the buyer may be uncertain as to the seller's beliefs (α). We assume that both buyer of type h and buyer of type l beliefs about the seller's beliefs are commonly known.[3] If the possible types of the seller are $\alpha^1 > \alpha^2 > ... > \alpha^n$, we will denote by β the probability that buyer type h assigns to seller type α^1 (the most optimistic of the buyer types). We assume without loss of generality that $\beta > 0$, since if this is not the case we can consider the most optimistic seller type that h considers with positive probability, and ignore more optimistic types altogether.

For example consider the following information structure[4]:

Seller \ Buyer	l	h
α^1	$0, 0$	$1, \frac{1}{4}$
α^2	$\frac{1}{3}, \frac{3}{8}$	$\frac{2}{3}, \frac{3}{8}$
α^3	$\frac{2}{3}, \frac{3}{8}$	$\frac{1}{3}, \frac{3}{8}$
α^4	$1, \frac{1}{4}$	$0, 0$

Here the seller is one of four types: she could either be certain of the buyer's valuation – types α^1 and α^4, or she could believe that one valuation is twice as likely as the other – types α^2 and α^3. The buyer's beliefs as to the seller's type assign a probability of $1/4$ that the seller actually knows the buyer's true type, and equal probability to the two uninformed types. We note that – as in this example – we do not require a common prior for the information structure[5]. Unlike the case of two-sided uncertainty about fundamentals, as

[3] See section 5 for a discussion of this assumption.

[4] The entries in the matrix specify the beliefs of the row (column) type of the seller (buyer) over the types of the buyer (seller). For example, in row 2 column 1 ($\frac{1}{3}, \frac{3}{8}$) stand for the seller type α^2 believes that the buyer has type l with probability $\frac{1}{3}$, and the buyer of type l believes that the seller is of type α^2 with probability $\frac{3}{8}$.

[5] There is no common prior since if assume a common prior we have that if the pair of types (α^2, l) prior is x then by the seller's posteriors the prior of (α^2, h) is $2x$, by the buyer posteriors this implies that the prior for (α^3, h) is $2x$ and by the seller we get the

in Yildiz (2001) where disagreement (lack of a common prior) over who gets to make the next offer is shown to have dramatic results on the outcome, relaxing the common prior assumption has no impact on the results in our case.

A distinction should be made between a seller who assigns probability one to a buyer's true type and a seller who *knows* the buyer's type. For our results the latter case coincides with the case where the type of seller who assigns probability one, for example to h, himself is assigned probability zero by the *other* buyer type, i.e. by l. Hence the information structure stated above also represents the case where the seller has types that *know* the buyer's type. Finally, we assume that the details of the information structures as stated above are commonly known.

The bargaining game between the seller and the buyer is defined as follows. Each time t the seller makes an offer p_t and the buyer either accepts or rejects the offer at that time. If the buyer never accepts an offer, the payoffs are zero to both players. If the buyer accepts an offer p_t at time t then the payoff to the seller is p_t discounted to time t and the payoff to the buyer is his valuation (l or h) minus p_t discounted to time t with both players having the same discount factor δ.

3 Uncertainty about Uncertainty and Delay

Consider a two-sided private information structure as described in (1) with parameters $0 < \alpha, \beta, \gamma < 1$. We have a buyer with private valuation and a seller that could be informed or uncertain (uninformed) about this valuation. There is common knowledge of gains from trade. Also, since it is assumed that the seller cannot be certain of h (resp. l) when the buyer is actually l (resp. h), we find that in this case the players cannot agree to disagree about the size of the gains from trade, i.e., we always have a common prior[6].

prior for (α^3, l) is $4x$, using the buyer posteriors we get the prior for (α^2, l) is $4x$, since $x \neq 0$ (otherwise all states would have a prior of 0) contradicting a common prior.

[6] As we mentioned earlier a common prior is not required for our results; rather we wish to emphasize that they are not driven by a lack of a common prior.

Seller \ Buyer	l	h
I^h	$0, 0$	$1, 1-\beta$
U	$1-\alpha, \gamma$	α, β
I^l	$1, 1-\gamma$	$0, 0$

(1)

We will show that for any given $0 < \alpha, \beta < 1$, delay does not vanish as the seller is allowed to make frequent offers. More precisely, we will show that as the discount factor δ goes to 1, each corresponding bargaining game has a Bayesian perfect equilibrium such that the delay in agreement increases in direct proportion to the frequency of offers made. We provide a precise description of the distribution of the time until agreement in equilibrium.

The equilibrium we construct is within the class of sequential equilibria satisfying the properties OP and RD defined below:

Definition 1 *A sequential equilibrium is said to satisfy the* optimistic pure strategy property (OP) *if the seller type that assigns the highest probability to the high valuation buyer is playing a pure strategy.*

Equilibria satisfying OP are similar to the equilibria constructed in the reputation literature. This definition mirrors the equilibria used in Kreps and Wilson (1982), Milgrom and Roberts (1982) and Abreu and Gul (2000), where the strong type follows a pure (and even stationary) strategy and the other types try to build a reputation of being strong by mimicking that behavior – mixing and increasing the posterior probability of being the strong type. In our model all types are fully rational, and the "strong" type is simply a type that assigns a higher probability to h.

Definition 2 *A sequential equilibrium is said to satisfy the* revealing deviation property (RD) *if off the equilibrium the buyer assigns zero probability to the most optimistic type.*

Without loss of generality we assume that informed types *know* the buyer's type. Hence, type h always assigns probability 0 to I^l on and off the equilibrium path. Therefore, RD implies that a deviation leads h to assign probability 1 to U in the case depicted in (1).

Once again the equilibria presented in the papers mentioned above follow this property. The equilibrium we construct when we show that a delay can occur satisfies even stronger properties. Our construction has the optimistic player not only following a pure strategy but actually offering a non-increasing price or, to be more precise, a constant price until the point at which he is revealed. This property can be seen as a form of stationarity and seems to follow the definition in Cho (1990) and the stronger notion in Fudenberg and Tirole (1991 p.408)[7].

We point this out to emphasize that it is not the lack of stationarity that drives our result[8]. What we show is not only that one can get delay within the class of sequential equilibria satisfying OP and RD, but also that with two-sided (first order) private information it is *impossible* to get delay within this class. Hence we are capturing properties distinct to higher order beliefs.

We are now ready to state the main result:

Theorem 1 *For every* $0 < \alpha, \beta < 1$ *and price* P *with* $l < P < h$ *for any* δ *close enough to* 1 *one can find a sequential equilibrium* $\zeta(\delta)$ *that satisfies* OP *and* RD *and such that the number of times that* $p_t = P$ *occurs with positive probability according to* $\zeta(\delta)$ *goes to infinity as* $\delta \to 1$*. Moreover, if we let* $\delta = e^{-r\Delta}$*, with* Δ *being the real time between two offers, then as* $\Delta \to 0$ *(i.e. frictions disappear) the expected delay (the expectation of time* $T\Delta$ *such that* $T = \max\{t | p_t = P\}$*) is bounded away from zero.*

The second part of the theorem assures that the delay we observe is indeed substantial. If we consider $\xi(\delta)$ to be the random variable measuring the number of periods t until the bargaining game stops according to $\zeta(\delta)$, the theorem implies that $E(\xi(\delta)) \underset{\delta \to 1}{\longrightarrow} \infty$.

Below we provide the proof of the theorem for the case where $\alpha < l/h$. The general case is proven in the appendix. The reason why the case $\alpha < l/h$ is relatively easier is that with such initial beliefs, whenever the seller might be revealed to be uninformed, the unique Bayesian perfect equilibrium in the corresponding continuation of the game is for her to offer price l immediately. When $\alpha \geq l/h$ we may be in a situation in which the seller is revealed to be uninformed but still assigns a high probability to the buyer being of type

[7] This property can also be seen as analogous to a monotonicity property when a continuum of types is present. Replacing our two types of sellers with a continuum of sellers of each type, where each seller plays a pure strategy, the condition OP translates to the existence of a cut-off point in the collection of less optimistic types. All types above the cut-off play according to the optimistic types' pure strategy and all below play another pure strategy. This construction has all the optimistic types play a pure strategy and other types are split between mimicking that pure strategy and playing another pure strategy.

[8] We would like to thank In-Koo Cho for pointing out the relation to the stationarity conditions in bargaining games.

h. This will lead to the seller offering a sequence of prices that decrease to l very quickly while the buyer is randomizing. We abuse the notion of a subgame and call the continuation of the game once the type is revealed a "subgame". This is justified by the fact that the buyer's beliefs at that point are singletons. Hence the players are practically playing in a subgame. This subgame corresponds to the case of one-sided information, hence *at the subgame* the Coase conjecture holds. The difficulty with this case is mostly technical since once types are revealed the actual subgame being played can depend on the number of offers that were made, and hence introduces a more complicated backtracking of future payoffs.

Proof. Let α be such that $0 < \alpha < l/h$, assume P is between l and h and let $0 < \beta < 1$. Consider the following (partially described) strategies. The strong types are I^h and l. The randomizing types (trying to build reputation) are U and h.

Type I^l offers l at each period t, no matter what the history is. Type I^h offers P at every period when the buyer's beliefs put probability less than 1 on this type (otherwise she offers $p = h$), i.e., this type offers P at the beginning and keeps offering P as long as she believes that the buyer assigns positive probability to type U where these beliefs are determined by the mixed strategy of type U described below. Type l only accepts an offer l or lower.

Type U chooses to offer l or P at time t with probabilities σ_t and $1 - \sigma_t$ respectively, and type h always accepts l or lower offers and accepts the offer P at time t with probability μ_t. Behavior off the equilibrium path (for prices above l other than P) will be described shortly.

Given this behavior we have the following beliefs of the types U and h along the equilibrium path:

α_t = the probability that the uninformed seller U assigns to the buyer being of type h at time t (before the buyer responds to the offer at time t and after all previous offers were P and were rejected).

β_t = the probability that buyer h assigns to the seller being of the uninformed type at time t (before the seller makes an offer at time t and given all previous offers were P and were rejected)[9].

The first-period beliefs are $\alpha_1 = \alpha$ and $\beta_1 = \beta$.

[9]The stationarity of this equilibrium is with respect to these state variables – α_t and β_t.

Hence we have the following equations according to Bayes rule:

$$\alpha_{t+1} = \frac{\alpha_t(1-\mu_t)}{\alpha_t(1-\mu_t)+(1-\alpha_t)} \tag{2}$$

$$\beta_{t+1} = \frac{\beta_t(1-\sigma_t)}{\beta_t(1-\sigma_t)+(1-\beta_t)} \tag{3}$$

By choosing $\alpha_1 < l/h$ we get that in a sequential equilibrium type U must offer l immediately if she is revealed. This follows from the same argument as used in Fudenberg and Tirole (1991, section 10.2.5) for one-sided information case. When only the buyer has private information, the seller can never expect any payoff higher than a "take it or leave it" offer could yield. If the seller assigns a probability below l/h to the buyer being of type h, an offer above l could extract at most h with a probability below l/h, yielding a payoff lower than l. Hence there is no "take it or leave it" offer above l that would not be dominated by offering l and getting it for sure. This observation dramatically simplifies the situation at hand. Since α_t is non-decreasing, at a point of time that U is revealed she immediately offers l and the offer is accepted immediately, thereby terminating the game. This scenario demonstrates the similarities with the Kreps and Wilson (1982) model of two-sided private information. We note that the revelation of h can only occur *after* h accepts the price P and the game terminates. As to deviations by the seller to a price other than P yet above l, we assume that such offers are interpreted as revelation of the seller's U type – the RD property – and hence are rejected in that "sub-game" and followed by an l offer (given these beliefs it is the unique perfect Bayesian equilibrium in that "sub-game" as follows from Fudenberg and Tirole 1991)[10].

We now need to construct the exact mixing strategies that U and h use, show that these constitute a sequential equilibrium, and analyze how the probability of termination of the game behaves as the discount factor goes to one.

If the seller is mixing at time t she must be indifferent between mimicking I^h and deviating. Since once she deviates she will be revealed as the uninformed type, she expects her payoff to follow the game with one-sided information given a deviation.

[10] We recall that since the buyer's beliefs off the equilibrium in this case assign probability one to the uninformed type, even if it was the informed type that deviated, the beliefs of the buyer will not change from this point onward.

Her payoff today if she offers P is $\alpha_t\mu_t P + \delta(1-\alpha_t\mu_t)l$ whenever she expects either to randomize at time $t+1$ or to deviate to l for sure (in both cases the expected payoff at time $t+1$ is l since if she is mixing in equilibrium she must be indifferent). She would not strictly prefer to choose P at $t+1$, since that would make choosing P at time t strictly dominating; this follows from the property that the probability of P being accepted is non-increasing. But she is randomizing at time t which implies that the payoff she expects from offering P is equal to the payoff from offering l and we have

$$\alpha_t\mu_t P + \delta(1-\alpha_t\mu_t)l = l \tag{4}$$

Rewriting (4) we have

$$\alpha_t\mu_t = \frac{l-\delta l}{P-\delta l} \tag{5}$$

Using (2) and (5) we get

$$\alpha_t = (1-\frac{l-\delta l}{P-\delta l})\alpha_{t+1} + \frac{l-\delta l}{P-\delta l} \tag{6}$$

Denote by α^* the minimal probability such that the seller might consider mixing, i.e. for all $\alpha < \alpha^*$, even if the buyer of type h will accept P for sure, the seller U strictly prefers to offer l:

$$\alpha^* = \frac{l-\delta l}{P-\delta l} \tag{7}$$

The crucial property of (7) is that as δ goes to 1 we have α^* going to 0 at the same rate (in particular for large enough δ we have α^* below α_1).

Similarly, if the buyer is randomizing, he is indifferent between accepting the offer P at time t or rejecting and getting the expected payoff at time $t+1$. The payoff at time $t+1$ is $h-P$ if the seller offers P, or the payoff will be $h-l$ if the seller reveals herself to be uninformed at time $t+1$. Hence we have

$$h-P = \delta((1-\beta_{t+1}\sigma_{t+1})(h-P) + \beta_{t+1}\sigma_{t+1}(h-l)) \tag{8}$$

or

$$\beta_t\sigma_t = \frac{(1-\delta)(h-P)}{\delta(P-l)} \qquad \text{for } t>1 \tag{9}$$

using (3) and (9) we have

$$\beta_t = (1-\frac{(1-\delta)(h-P)}{\delta(P-l)})\beta_{t+1} + \frac{(1-\delta)(h-P)}{\delta(P-l)} \tag{10}$$

Let

$$\beta^* = \frac{(1-\delta)(h-P)}{\delta(P-l)} \tag{11}$$

For every $\beta < \beta^*$ the buyer will be better off accepting P, and for every $\beta > \beta^*$ there is a probability $\sigma_t < 1$ such that the buyer would not mind waiting another period.

Consider any α_1, β_1 such that $0 < \alpha_1 < l/h$ and $0 < \beta_1 < 1$. Let δ be close enough to 1 such that $\alpha^* < \alpha_1$ and $\beta^* < \beta_1$. Consider the sequence $\alpha(1) = \alpha^*, \alpha(2) = (1-\alpha^*)\alpha^* + \alpha^*, \alpha(3) = (1-\alpha^*)\alpha^2 + \alpha^*, ...$ this sequence follows backtracking (6) as is shown in Figure 1. Solving this recursion yields

$$\alpha(n) = (1-\alpha^*)^{n-1}\alpha^* + (1-\alpha^*)^{n-2}\alpha^* + ... + (1-\alpha^*)^0\alpha^*$$

or

$$\alpha(n) = 1 - (1-\alpha^*)^n$$

let $N(\alpha_1) = \max\{n|\alpha(n) < \alpha_1\}$. Similarly we define $M(\beta_1) = \max\{n|\beta(n) < \beta_1\}$ where $\beta(n) = 1 - (1-\beta^*)^n$ is generated by iterating (10). Since α_1, β_1 are given, we denote $N = N(\alpha_1)$ and $M = M(\beta_1)$.

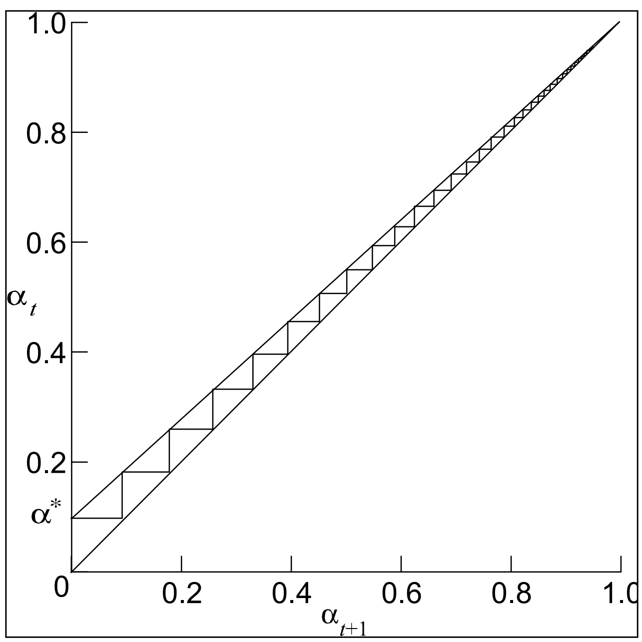

Iterations of the equations $\alpha_t = (1-\alpha^*)\alpha_{t+1} + \alpha^*$

Figure 1

We distinguish two possible cases

1. $N \leq M$
2. $N > M$

In the first case we define the following strategies: Let σ_1 be such that $\beta_2 = \beta(N)$. This can be obtained by (3) since

$$\beta_1 > \beta(M) \geq \beta(N) \tag{12}$$

for all $2 \leq t \leq N+1$ let σ_t be such that (9) is satisfied which implies that $\beta_t = \beta(N+2-t)$ and in particular $\beta_{N+1} = \beta^*$ and $\sigma_{N+1} = 1$. For $1 \leq t \leq N+1$ let α_t be defined by (6) and let μ_t be such that (5) is satisfied for $1 \leq t \leq N$, let $\mu_{N+1} = 1$. By definition the beliefs α_t, β_t follow Bayes rule for $t = 1, ..., N+1$ according to the strategies stated above. For $t = 1, ..., N$ both players are indifferent between the two choices over which they are randomizing, since (6) and (10) are satisfied. We need only show that at $t = N+1$ both players are playing a best response and that the continuation beyond this period (if any) is well defined. But at $t = N+1$ the uninformed seller is choosing l with probability 1 just after the buyer's beliefs have reached β^* and so the buyer will definitely accept P at $t = N+1$ since he will deduce that the seller type must be I^h and we have that $\mu_{N+1} = 1$ is a best response. As for the buyer, since $\alpha_{N+1} < \alpha^*$, the seller strictly prefers to jump to l and receive l for sure at that period, from either type, rather than receive P for sure if the buyer is the h type or δl if it is the l type since $\alpha_{N+1}P + \delta(1-\alpha_{N+1})l < l$. This implies that $\sigma_{N+1} = 1$ is a best response as required. These strategies have the game terminate at most at period $N+1$. Furthermore, with positive probability the game lasts N periods.

For case 2 ($N > M$) we construct the following strategies. Let $\sigma_1 = 0$ and σ_t satisfy (9) for all $2 \leq t \leq M+1$. Let μ_1 be such that $\alpha_2 = \alpha(M)$ as derived from (2). Since

$$\alpha_1 \geq \alpha(N) > \alpha(M) \tag{13}$$

(note the strict inequality), we have that $\alpha_1 \mu_1 > \frac{l-\delta l}{P-\delta l}$ which implies that the seller choice of P in the first offer is a strict best response. For $2 \leq t \leq M$ let μ_t follow (5). At period $t = M$ the buyer is still mixing since $\beta_M > \beta^*$ and this leads to the seller beliefs being $\alpha_{M+1} = \alpha^*$. Since $\beta_{M+1} < \beta^*$, the buyer will accept P at time $M+1$ which is exactly why a seller at $\alpha_{M+1} = \alpha^*$ will be indifferent and can follow σ_{M+1} as required.

We have just shown that in the constructed equilibrium there is positive probability that the game will continue at least $Min\{M(\delta), N(\delta)\}$ periods before agreement is reached. We have that $\alpha_1 \approx 1-(1-\alpha^*)^{N(\delta)}$ and $\beta_1 \approx 1-(1-\beta^*)^{M(\delta)}$ but from the definition of α^* and β^* we have that $\alpha^* \approx (1-\delta)\frac{l}{P-l}$ and $\beta^* \approx (1-\delta)\frac{h-P}{P-l}$ and so $N(\delta) \approx \ln(1-\alpha_1)/\ln(1-(1-\delta)\frac{l}{P-l})$ and $M(\delta) \approx \ln(1-\beta_1)/\ln(1-(1-\delta)\frac{h-P}{P-l})$. If we fix the discount rate at λ and let the frequency of offers increase[11], then the number of offers within a time period behaves according to $-1/\ln\delta$ but $\ln(1-(1-\delta)K) \approx \ln\delta$ as $\delta \nearrow 1$ for a positive constant K. We can conclude that the maximal length of time the players do not trade with positive probability grows at the *exact* same rate as the frequency of offers grows within a given time period.

It is readily seen that our construction yields a sequential equilibrium satisfying OP and RD by construction. Finally, let's consider the stationarity of the constructed equilibrium. Along the equilibrium path, after type U is revealed the unique equilibrium in the continuation subgame is stationary. Before type U is revealed the mixed behavior strategies depend only on the current beliefs, as described by equations (5), (9) and the conditions for mixing in the first round. Off the equilibrium path the strategies depend only on the current beliefs. So this equilibrium also satisfies stationarity in the spirit of Cho (1990) and the stronger notion in Fudenberg and Tirole (1991 p.408). ■

We have shown that with positive probability, agreement is not reached until the number of possible offers above l (and bounded away from l) is of the order of $-1/\ln\delta$. This violates the Coase conjecture but in a weak sense. After all, the probability that agreement is not reached up until that last possible offer above l is going to zero as δ goes to one (although the time period of that last offer grows to infinity). It turns out that a stronger result holds. Namely, the expected time until agreement is reached is bounded away from zero as δ goes to one. In the next subsection, we will turn to the continuous time approximations of equilibrium behavior for the study of the distribution of delay time and total excepted payoffs in our bargaining game, but first we discuss equilibria with delay in more general cases.

General structure of seller's types.

When we described the model we provided an example of an information

[11] For example consider $\delta_n = \lambda^{1/n}$, which corresponds to making n offers at every fixed time period with $\lambda = e^{-r}$. Hence $n = \ln\lambda/\ln\delta_n$.

structure with 4 types for the seller (these correspond to I^h, two uninformed types with different α, and I^l). Up until now we have constructed equilibria only in the simplest case of one uninformed type. It turns out that the equilibria in the more general case can be easily constructed along the same line. Consider a setup with two informed types I^h and I^l and m uninformed types with prior beliefs α^i $i = 1, ..., m$ ranked by how optimistic they are: $I^h = \alpha^0 > \alpha^1 > \alpha^2 > ... > \alpha^m > \alpha^{m+1} = I^l$. First notice that if a type α^i is mixing in a given period t in equilibrium, then all more optimistic types are strictly better off offering P under the OP condition and all more pessimistic types are strictly better off revealing themselves to be uninformed. To construct the equilibrium, we start from time T when the most optimistic seller exits. Roll back the equations, as in the equilibrium we constructed for the one type case, as if there are no other types until we reach the prior beliefs of one of the players (buyer or seller α^1). If it is the seller mixing in the first round, then the equilibrium with many types is that all the sellers type $\alpha^2, ..., \alpha^m$ follow the one-sided uninformed strategy at the first round revealing themselves to be uninformed (for small enough values of α they simply offer l) and the seller α^1 follows the equilibrium strategy described above. If it is the buyer that mixes in the first round, then we keep backtracking with seller α^2, now we roll back (from the point we reached in the previous roll-back) assuming that seller type α^2 and the buyer are randomizing and that the seller α^1 offers P with probability 1 like the informed seller. We roll back until we reach one of the priors: if we reach the seller α^2 first then we are done; if it is the buyer, then we consider α^3 and so on. The behavior generated by this procedure has exactly one mixing uninformed seller type at any given period (other than, perhaps, the first period) and all other types are pooling with the informed types. At a given period the more optimistic sellers pool with I^h and keep offering P while the more pessimistic types reveal that they are uninformed and follow the one-sided information behavior: making lower offers going quickly to l. Constructing off the equilibrium behavior according to RD yields the required generalization. However, we need to say how RD translates to this general case, since RD states that a deviation is assumed to be made by the *single* uninformed type. We can assume that in the general case RD translates to the buyer believing that a deviation was made by the most pessimistic type and the proof follows easily. We note that RD could be relaxed in the general case by assuming that the buyer believes the deviation *could not* have been made by the informed type, i.e. his beliefs off the equilibrium are supported on the uninformed types of the seller. By excluding the informed type we will be in the completely mixing seller type case when a deviation occurs. This assures fast convergence to l according to our results in Section 4; however, it complicates the backtracking of the equilibria behavior, hence the detailed proof is omitted.

So the equilibria constructed in this section are robust to the introduction of more uninformed types. As we show in Section 4, the critical element for the existence of equilibria with delay is the possibility of types that do not have fully mixed beliefs – the informed types.

3.1 Continuous Time Approximations

Before we begin studying the approximations of equilibrium strategies presented above, it is important to emphasize that we do not analyze a continuous time game, but rather the continuous limit of the equilibrium play path. We are not aware of a known definition of a continuous time game that allows for strategies that are the continuous limit of our discrete time equilibrium strategies. The reason for this is that we would like to have strategies where both players are continuously mixing. Hence we cannot use the formulation of games where players are required to commit to a pure strategy for an arbitrary (small) period of time. We also cannot use differential game forms since we do want to allow discontinuous price choices. It is worthwhile noting the difference between the general bargaining games with two-sided private information that we consider, and the construction in Kreps and Wilson (1982) and Abreu and Gul (2000). The latter avoid the difficulties of the continuous time formulation since once a 'jump' is made, the game terminates and thus the continuous time game is well defined as a game of choosing a stopping time. Since this approach cannot be utilized in our framework, nor in general two-sided private information models, we consider the sequence of equilibria that we construct and study their behavior as δ goes to 1. This turns out to be fairly straightforward since the equations determining the mixing strategies have a continuous time formulation. As δ goes to 1 we can replace the decrease in the discounting of the next offer with a constant discount but with offers being made more and more frequent. The distribution of the time until agreement is reached as determined by the discrete equations converges to the distribution generated by the continuous time equations obtained in this way.

Denote by r the discount rate that the players use for a given unit of time. Denote the original beliefs by α_0 and β_0 respectively. Denote by dt the time between offers. Then $\delta = e^{-rdt}$. Redefine the probabilities to be "per time between offers," so that the seller mixes with probability $\sigma_t dt$ and the buyer mixes with probability $\mu_t dt$.

With this formulation, first notice that continuous time analogs of the Bayes rules (2) and (3) are:

$$\dot{\alpha}_t = -\alpha_t(1-\alpha_t)\mu_t \tag{14}$$

$$\dot{\beta}_t = -\beta_t(1-\beta_t)\sigma_t \tag{15}$$

Second, become after taking the limit $dt \to 0$, equations (5) and (9):

$$\alpha_t \mu_t = r\left(\frac{l}{P-l}\right) \tag{16}$$

$$\beta_t \sigma_t = r\left(\frac{h-P}{P-l}\right) \tag{17}$$

These four equations are the continuous time approximations of the equilibrium conditions specified in Theorem 1 [12]. Equations (2), (3), (5) and (9) together with definitions of α^*, β^* and the probabilities in the first round of the game fully characterize equilibrium, so we can use the continuous their time limits as an approximation of how the equilibrium looks for a game with frequent offers. Note that both β^* and α^* converge to 0, so in the limit the posteriors α_t and β_t converge to 0 at the same time.

Denote $B = r\left(\frac{h-P}{P-l}\right)$ and $A = r\left(\frac{l}{P-l}\right)$. Using equations (16) and (14) we get:

$$\dot{\alpha}_t = -A(1-\alpha_t) \tag{18}$$

and using equations (17) and (15) we get:

$$\dot{\beta}_t = -B(1-\beta_t) \tag{19}$$

Integrating yields the paths for the two beliefs along the equilibrium path:

$$\beta_t = 1-(1-\beta_{0+})e^{Bt} \tag{20}$$

$$\alpha_t = 1-(1-\alpha_{0+})e^{At} \tag{21}$$

[12] These equations are approximations of the best responses also in the general case $\alpha_0 > \frac{l}{h}$, because as $\delta \to 1$ after the seller reveals himself to be uninformed, then the first price he offers converges to l (the Coase conjecture for the limit behavior in that 'subgame').

where β_{0+} and α_{0+} denote the beliefs at time 0 after the possible atom at time 0. The atom corresponds to the behavior in the first round of the game. Dividing (19) and (18) and integrating we obtain $(1-\beta_t) = k(1-\alpha_t)^{B/A}$. Since α_t and β_t converge to 0 at the same moment, we have $k = 1$. That pins down the beliefs at time $t = 0^+$ (as well as the whole path):

$$1 - \beta_{0+} = (1 - \alpha_{0+})^{B/A} \tag{22}$$

If the original beliefs β_0 and α_0 do not satisfy this condition, then one of the players 'exits' with an atom at time 0, as we described in the proof of Theorem 1. If

$$1 - \beta_0 < (1 - \alpha_0)^{B/A} \tag{23}$$

then in equilibrium the uninformed seller offers l with an atom at time 0. If the inequality goes the other way, the h buyer accepts P with an atom at time 0. The direction of this inequality depends on the original beliefs and P (as $B/A = \frac{h-P}{l}$). Note that the game surely ends by time $T^* = \min\{\frac{-\ln(1-\alpha_o)}{A}, \frac{-\ln(1-\beta_o)}{B}\}$. For any $\beta_0 \in (0,1)$ and $\alpha_0 \in (0,1)$ the atom at time 0 is bounded away from 1, together with (22), (20) and (21), we have a strictly positive probability that the game will not end instantly so indeed the expected length of the game is strictly positive, as claimed.

3.2 Expected Payoffs and Delay

With the continuous time approximations of the equilibrium strategies constructed above, we can calculate the players' expected payoffs and expected delay. They depend crucially on who has an atom at time 0, which in turn is determined by how P in the equilibrium we construct relates to α_0 and β_0. Using equations (22) and (23) we can show that the critical value of P is:

$$P^* = h - l\frac{\ln(1-\beta_0)}{\ln(1-\alpha_0)} \tag{24}$$

For $P > P^*$ the U seller offers l with an atom at time 0, and for $P < P^*$ the h buyer accepts P with an atom at time 0. We first look at a case when the game does not end at time 0. The hazard rates σ_t and μ_t yield the distribution of stopping times by buyer and seller:

$$f_S(t) = \sigma_t \exp(-\int_0^t \sigma_\tau d\tau) \tag{25}$$
$$f_B(t) = \mu_t \exp(-\int_0^t \mu_\tau d\tau)$$

Using equations (16), (17), (20) and (21) we get:

$$f_S(t) = \frac{B}{\beta_{0^+} e^{Bt}} \text{ for } t \in (0, T^*] \tag{26}$$
$$f_B(t) = \frac{A}{\alpha_{0^+} e^{At}} \text{ for } t \in (0, T^*]$$

It turns out that the calculation of profits of the relevant types uses only $f_B(t)$. This is so, since conditional on the game not ending at time 0, both the h buyer and the U seller are indifferent between immediate termination and delay for another dt. So their expected profits are[13]:

$$E\left(\Pi_h|(t = 0^+)\right) = h - P \tag{27}$$
$$E\left(\Pi_U|(t = 0^+)\right) = l \tag{28}$$

The payoff of the I^h seller is given by:

$$E\left(\Pi_{I^h}|(t = 0^+)\right) = \int_0^{T^*} Pe^{-rt} \frac{A}{\alpha_{0^+} e^{At}} dt$$

Integrating and using the definition of A we get:

$$E\left(\Pi_{I^h}|(t = 0^+)\right) = \frac{l}{\alpha_{0^+}} \left(1 - (1 - \alpha_{0^+})^{\frac{P}{l}}\right) \tag{29}$$

Total payoffs depend crucially on the strategies at time 0. If $P = P^*$, then no player exits with an atom and the total expected payoffs are given by the expressions above – there is no need to condition on $t = 0^+$ in this case.

[13] We talk only about three relevant types as trivially the payoff of the l buyer is 0 and the payoff of I^l seller is l.

If $P < P^*$ it is the h buyer that 'exits' with an atom. Given a prior α_0, atom μ_0 has to be:

$$\mu_0 = \frac{1}{\alpha_0} - \frac{1-\alpha_0}{\alpha_0}(1-\beta_0)^{-\frac{l}{h-P}} \tag{30}$$

to satisfy the Bayes rule and equation (22). If $P > P^*$ it is the U seller that exits with an atom. Given a prior β_0 the atom σ_0 has to be:

$$\sigma_0 = \frac{1}{\beta_0} - \frac{1-\beta_0}{\beta_0}(1-\alpha_0)^{-\frac{h-P}{l}} \tag{31}$$

Summing up all these calculations, the total expected payoff for the I^h seller is:

$$E\left(\Pi_{I^h}(P)\right) = \left\{ \begin{array}{c} \frac{l}{\alpha_0}\left(1-(1-\alpha_0)^{\frac{P}{l}}\right) \text{ for } h > P \geq P^* \\ \mu_0 P + (1-\mu_0)\frac{l}{\alpha_{0^+}}\left(1-(1-\alpha_{0^+})^{\frac{P}{l}}\right) \text{ for } l < P < P^* \end{array} \right\} \tag{32}$$

where $\alpha_{0^+} = 1-(1-\beta_0)^{\frac{l}{h-P}}$ from equation (22). Expected payoff for the uninformed seller is:

$$E\left(\Pi_U(P)\right) = \left\{ \begin{array}{c} l \text{ for } h > P \geq P^* \\ \alpha_0\mu_0 P + (1-\alpha_0\mu_0)l \text{ for } l < P < P^* \end{array} \right\} \tag{33}$$

and expected payoff for the h buyer:

$$E\left(\Pi_h(P)\right) = \left\{ \begin{array}{c} \beta_0\sigma_0(h-l) + (1-\beta_0\sigma_0)(h-P) \text{ for } h > P \geq P^* \\ h-P \text{ for } l < P < P^* \end{array} \right\} \tag{34}$$

Note that all the expected payoffs are independent of the discount rate r.

How do these expected payoffs vary with P? The payoff of the I^h seller is increasing in P in the first region and is in general non-monotonic in the second. So the maximum is attained either at $P = h$ or for some $P < P^*$. The payoff of the U seller is strictly concave for $P \leq P^*$. At $P = P^*$ and $P = l$ it is equal to l. So a local maximum exists, is unique and lies somewhere between these two values. The payoff of the h buyer is maximized for $P = l$ (the buyers ranking of these equilibria is trivial for $P < P^*$ and quite complicated for high prices).

We now turn to the expected delay until agreement is reached. Naturally, the distribution of delay time depends on the configuration of types actually bargaining. Let $T(P)$ denote the random stopping time (time until agreement) and for a given P let F be the cumulative distribution function. If the types are U and h the distribution (c.d.f.) of the stopping time is:

$$F(t) = 1 - (1 - F_B(t))(1 - F_S(t)) \tag{35}$$

where the densities of F_B and F_S are defined in (26). Hence the distribution of $T(P)$ can be calculated much like the minimum of two independent exponential distributions (the minimum is also exponential with a parameter that is the sum of parameters of the two exponential distribution). That gives expected delay:

$$E\left(T(P)\right) = \left\{ \begin{array}{c} (1-\sigma_0)\int_0^{T^*} \frac{t(A+B)}{\alpha_{0+}\beta_{0+}\exp((A+B)t)} dt \text{ for } h > P \geq P^* \\ (1-\mu_0)\int_0^{T^*} \frac{t(A+B)}{\alpha_{0+}\beta_{0+}\exp((A+B)t)} dt \text{ for } l < P < P^* \end{array} \right\} \tag{36}$$

Given that σ_0 and μ_0 are both smaller than 1 and that $T^* > 0$ we can see that the expected delay is positive. Since $T^* = \min\{\frac{-\ln(1-\alpha_o)}{A}, \frac{-\ln(1-\beta_o)}{B}\}$ we have:

$$T^* = \left\{ \begin{array}{c} \frac{-\ln(1-\alpha_o)}{A} \text{ for } h > P \geq P^* \\ \frac{-\ln(1-\beta_o)}{B} \text{ for } l < P < P^* \end{array} \right\}$$

We do not calculate here how the expected delay varies with the fundamentals, α_0, β_0 and P, but note these can be derived from the above expressions. If the types are U and l or I^h and h the expected delay can be calculated similarly, by setting $F(t) = F_S(t)$ and $F(t) = F_B(t)$ respectively for even longer expected delay.

The integral in the expression for $E\left(T(P)\right)$ can be simplified to:

$$E(T(P)|t > 0)) = \frac{1}{\alpha_{0+}\beta_{0+}} \frac{1 - (T^*(A+B)+1)e^{-T^*(A+B)}}{(A+B)} \tag{37}$$

Since

$$T^*(A+B) = \left\{ \begin{array}{c} -(1+\frac{h-P}{l})\ln(1-\alpha_o) \text{ for } h > P \geq P^* \\ -(1+\frac{l}{h-P})\ln(1-\beta_o) \text{ for } l < P < P^* \end{array} \right\}$$

We have:

$$E\left(T(P)\right) = \left\{ \begin{array}{c} (1-\sigma_0)\frac{1}{\alpha_{0+}\beta_{0+}} \frac{1-(1-(1+\frac{h-P}{l})\ln(1-\alpha_o))(1-\alpha_0)^{1+\frac{h-P}{l}}}{r(\frac{h}{P-l}-1)} \text{ for } h > P \geq P^* \\ (1-\mu_0)\frac{1}{\alpha_{0+}\beta_{0+}} \frac{1-(1-(1+\frac{l}{h-P})\ln(1-\beta_o))(1-\beta_0)^{1+\frac{l}{h-P}}}{r(\frac{h}{P-l}-1)} \text{ for } l < P < P^* \end{array} \right\} \tag{38}$$

With the explicit calculation of the expected delay in (38) we now see that the expected delay is proportional to $1/r$.

To provide an idea of the size of the delay generated by our equilibria we suggest the following parameterization: we normalize $l = 1$ and set $h = 1.5l = 1.5$. Furthermore, we suggest $\beta = 0.9$ so only with 10% probability is the seller informed. We pick $r = 10\%$ as the *annual* discount rate. That leaves us with two free parameters: α_0 and P. We pick $\alpha_0 \in \{0.2, 0.5, 0.8\}$ and vary P between l and h. Figure 2 shows the expected delay for these parameters. The (unconditional) expected length of negotiations is of the order of $2 - 13$ *weeks* for $P = \frac{h+l}{2} = 1.25$.

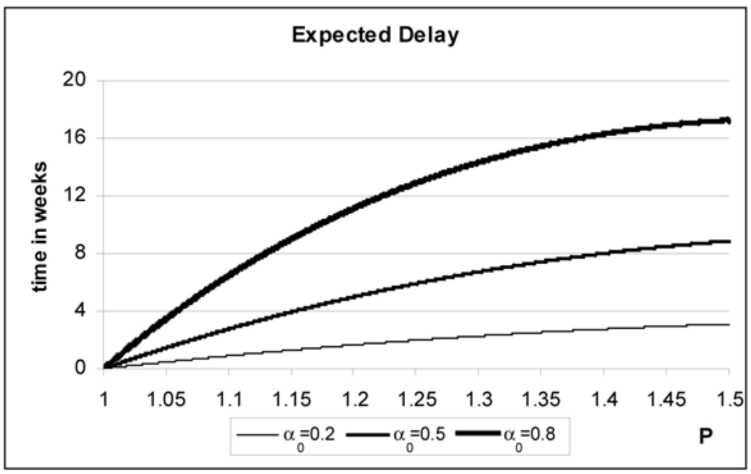

Figure 2

Given that in our parameterization the probability of the seller being informed is low, $P^* < l$, the uninformed seller is exiting with an atom at time 0 for all prices $l < P < h$. Figure 3 shows how this probability of no delay at all – by the uninformed type mixing at time 0 – changes with P. We see that, given that the seller is rarely informed, the probability that the uninformed seller tries to mimic him is quite small.

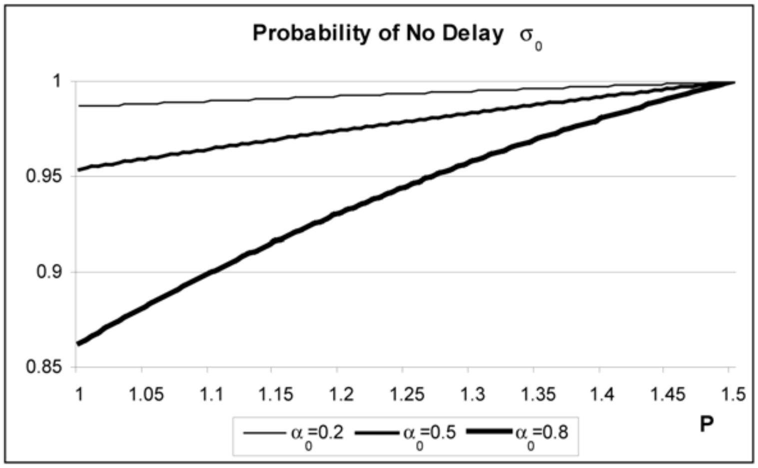

Figure 3

Finally, the expected delay conditional on the game not ending at time 0 can be calculated as $\frac{ET(P)}{(1-\sigma_0)}$. Given how large σ_0 is in our examples, the conditional delay is quite extensive. For example for $P = \frac{h+l}{2} = 1.25$ and our parameter values it turns out to be of the order of more than 4 years.

4 Completely mixed beliefs and no delay

In the previous section we have shown that uncertainty about uncertainty can dramatically change the equilibrium outcomes - in particular, introduce significant delay in equilibrium. In this section we show that a crucial component of the setup above is *possible exclusion of types:* the seller has a type with beliefs that put probability 0 on the buyer having low valuation. In contrast, in this section we consider a setup with the seller having two possible types (optimistic and pessimistic), both with fully mixed beliefs. We show that any equilibrium satisfying OP and RD (like the one constructed in Theorem 1) converges to the seller offering l immediately as $\delta \to 1$ (i.e. as offers become more and more frequent).

There are two types of seller: 1 and 2 (optimistic and pessimistic) and two types of buyers (as before). The prior belief structure is:

Seller\Buyer	l	h
1	$1-\alpha^1, \gamma$	$\alpha^1, 1-\beta$
2	$1-\alpha^2, 1-\gamma$	α^2, β

(39)

with:

$$1 > \alpha^1 > \frac{l}{h} > \alpha^2 > 0 \tag{40}$$

The remaining elements of the game are unchanged. We show that for given values α^1 and α^2 there exists a uniform (for all δ) upper bound on the number of periods the buyer rejects offers with positive probability in equilibrium and that the prices in all those equilibria converge to l for every round.

In the equilibria considered, the optimistic seller is the strong type and for high enough values of δ there does not exist an equilibrium with a constant price P offered by that optimistic type. To see this, consider an equilibrium of the form constructed in Theorem 1. We know that in finite time T, the pessimistic seller will reveal herself. After that we have the standard subgame, with only one type of seller. In that subgame the offered price converges to l as δ goes to 1. Hence, for high enough δ the buyer should not randomize the last few periods before T if the price offered at the end of those T periods is bounded away from l – so he is not playing a best response. That implies that the only candidate for an equilibrium with delay is where the equilibrium strategy for the optimistic seller is to follow a sequence of prices that converges to l.

Therefore the equilibria that we consider have the following structure: the optimistic seller follows a path of prices $\{P_t\}_0^\infty$ (a pure strategy by OP). The pessimistic seller randomizes between following this path and offering l (with probability $\sigma_t > 0$). The buyer randomizes between accepting the current price $P_t > l$ (with probability $\mu_t > 0$) and rejecting it. If any price different than P_t is ever observed, the buyer believes that the seller is pessimistic by the RD property.

Theorem 2 *With completely mixed types, in any equilibrium satisfying OP and RD the bargaining ends in at most N periods, where N is a uniform bound for all $\delta < 1$. Moreover, the price paths offered in these equilibria converge to l as offers are made more frequently.*

Proof. 1. By subgame perfection, in any equilibrium, prices that are accepted are in the range $[l, h]$ and beliefs α_t^i are weakly decreasing.

2. Consider any subgame in which the type of the seller is revealed at time T. If seller is pessimistic (type 2) the game ends immediately after T with the seller offering $p = l$. If the seller is optimistic (type 1) the equilibrium

strategy is uniquely determined. The prices follow a path:

$$p_{T+t} = h - \delta^{n(\alpha_T^1,\delta)-1-t}(h-l) \tag{41}$$

where $n(\alpha_T^1,\delta)$ is uniformly bounded for all δ given any $\alpha^1 \geq \alpha_T^1$. This is the equilibrium path described in Fudenberg and Tirole (1991). The bargaining takes at most $\bar{n}$ further rounds after that moment, with $\bar{n}$ uniformly bounded in δ.

3. When

$$\alpha_t^2 < \frac{l(1-\delta)}{h-\delta l} \tag{42}$$

the pessimistic seller offers l immediately (getting l for sure is better than getting h with probability α_t^2 and getting l next period with probability $(1-\alpha_t^2)$).

4. Consider any equilibrium in which the optimistic seller follows a pure strategy path $\{P_t\}_{t=0}^{\infty}$ and the beliefs out of equilibrium are such that any price different from that path reveals the seller to be type 2. Recall that these are the class of equilibria considered throughout the paper satisfying *PO* and *RD*. For any given δ there exists $T < \infty$ such that if the game is not over by period T then the pessimistic seller reveals herself at period T in equilibrium by offering l.[14]

The existence of such a T follows from the observation that had the pessimistic seller not revealed herself for sure at any period, then we would have for any t:

$$h\sum_{\tau=0}^{\infty} \delta^{\tau}\alpha_{t+\tau^t}^2 \mu_{t+\tau} \geq l \tag{43}$$

Given that $\alpha_{t+\tau}^2$ is nonincreasing this yields a lower bound on the probabilities of acceptance and hence in finite number of periods the posterior drops below $\frac{l(1-\delta)}{h-\delta l}$, making it optimal for the pessimistic seller to offer l for sure immediately.

So for any δ the pessimistic seller reveals herself with probability 1 at time T and the game lasts for at most $T + n(\alpha_T^1,\delta)$ periods (from now on we select T to be the earliest time that the seller reveals herself with probability 1).

[14] So the bargaining lasts for at most $T + n(\alpha_1^1,\delta)$ periods.

5. $n(\alpha_T^1, \delta)$ is bounded uniformly for all δ, so we only need to show that T is bounded uniformly for all δ. Consider a given δ and any period t before T: the buyer of type h is mixing so we must have:

$$h - P_t = \delta\left(\beta_{t+1}\sigma_{t+1}(h-l) + (1-\beta_{t+1}\sigma_{t+1})(P_t - P_{t+1})\right) \tag{44}$$

rearranging:

$$\beta_{t+1}\sigma_{t+1} = \frac{(h-P_t)(1-\delta) - \delta(P_t - P_{t+1})}{\delta(P_{t+1} - l)} \tag{45}$$

As $h > P_t > l$ and $\beta_{t+1}\sigma_{t+1} > 0$ that restricts the change in prices:

$$P_t - P_{t+1} < (h-l)\frac{1-\delta}{\delta} \tag{46}$$

It implies that for high δ the prices cannot drop too fast.

As at time T the pessimistic seller offers l and reveals herself, the optimistic seller has to offer

$$P_T = h - \delta^{n(\alpha_T^1,\delta)-1}(h-l) \tag{47}$$

and follow the path of the game with complete information about the seller, where $n(\alpha_T^1, \delta)$ is uniformly bounded for all δ for a given α_T^1 and weakly decreasing in α_T^1. So $\lim\limits_{\delta\to 1} P_T = l$. Combining (46) and (47) bounds the price in period $T - \tau$ to at most:

$$P_{T-\tau} < P_T + \tau(h-l)\frac{1-\delta}{\delta} \tag{48}$$

which for a given τ converges to l as δ goes to 1.

Now consider the pessimistic seller. He is mixing at every period before T so for $t < T$:

$$l = P_t\alpha_t^2\mu_t + \delta(1-\alpha_t^2\mu_t)l \tag{49}$$

which implies:

$$\mu_t\alpha_t^2 = \frac{(1-\delta)\,l}{(P_t - \delta l)} \tag{50}$$

Given $\alpha_t^2 < \frac{l}{h}$ we get:

$$\mu_t > \frac{(1-\delta)h}{(P_t - \delta l)} \tag{51}$$

In any period $T - \tau$ we have:

$$\mu_{T-\tau} > \frac{(1-\delta)h}{(h - \delta^{n(\alpha_T^1,\delta)-1}(h-l) + \tau(h-l)\frac{1-\delta}{\delta} - \delta l)} \tag{52}$$

As δ converges to 1 the right side converges to:

$$\frac{1}{((h-l)(\overline{n}-1+\tau)+l)\,h} \tag{53}$$

where $\overline{n}$ is the limit of $n(\alpha_T^1, \delta)$ as δ goes to 1.

So for any number of periods τ, we have that $\mu_{T-\tau}$ is bounded uniformly from zero. Finally, consider the threshold α^* such that if $\alpha_t^2 < \alpha^*$, the pessimistic seller prefers to obtain l immediately rather than $P_{T-\tau}$ with probability α_t^2 in the current period and l next period with probability $(1-\alpha_t^2)$. For α^* we get:

$$P_{T-\tau}\alpha^* + \delta(1-\alpha^*)l = l$$

$$\alpha^* = \frac{(1-\delta)l}{P_{T-\tau} - \delta l}$$

As δ converges to 1 we have:

$$\lim_{\delta\to 1} \alpha^* = \frac{1}{((h-l)(\overline{n}-1+\tau)+l)} > 0 \tag{54}$$

Combining these results, for any $\alpha^2 < \frac{l}{h}$ in a uniformly bounded number of steps the beliefs have to drop below α^* which implies a uniform upper bound on T.

6. We have established that the number of bargaining rounds is uniformly bounded (so bargaining ends quickly with frequent offers. Finally, as

$$P_1 \leq P_T + T(h-l)\frac{1-\delta}{\delta} \tag{55}$$

all prices ever offered in the equilibrium converge to l as well. ■

4.1 Continuous time approximation and no delay

We now look at a continuous time approximation for the case of completely mixed beliefs with higher order uncertainties. We use the same notation as before: original beliefs are α for the pessimistic seller's belief about the buyer being of high value, and β for the buyer's belief about the seller's type. As before with dt as the time between offers, we have $\delta = e^{-rdt}$. The seller mixes with probability $\sigma_t dt$ and the buyer mixes with probability $\mu_t dt$. The main difference now is that the price is not constant. Denote by P_t the continuous-time limit of the price path followed by the optimistic type in equilibrium.

Notice that any equilibrium in Theorem 2 can be described by 4 equations: two Bayes rules and equations (45) and (50). The continuous time limit of the Bayes rules are (14) and (15). Taking the limits as $dt \to 0$ of (45) and (50) yields:

$$\beta_t \sigma_t = \frac{\dot{P}_t + r(h - P_t)}{P_t - l} = B(t) \tag{56}$$

$$\alpha_t \mu_t = \frac{rl}{P_t - l} = A(t) \tag{57}$$

Using Bayes rule we obtain:

$$\alpha_t = 1 - (1 - \alpha_{0+}) e^{\int_0^t A(\tau) d\tau} \tag{58}$$

$$\beta_t = 1 - (1 - \beta_{0+}) e^{\int_0^t B(\tau) d\tau} \tag{59}$$

Using this approximation, we now show that delay is not possible with frequent offers. For the equilibrium to hold, it has to be the case that P_t converges to l at the same time that β_t converges to 0[15]. That has to happen in bounded time: If $\alpha_t = 0$ at any time, the dominant strategy for the pessimistic seller is to offer l immediately and hence the price has to converge at that moment to l. As $A(t) \geq \frac{rl}{h-l} > 0$ we get:

$$0 \leq \alpha_t \leq 1 - (1 - \alpha_{0+}) e^{t \frac{rl}{h-l}} \tag{60}$$

The right side is strictly decreasing in t and reaches 0 in finite time. So there exists T such that $\lim_{t \to T} P_t = l$. Suppose $T > 0$. At time T we need $\alpha_t \geq 0$ and $\beta_t \geq 0$. That requires $\int_0^T A(\tau) d\tau$ and $\int_0^T B(\tau) d\tau$ to be bounded. Denote

[15] In a subgame with only types 1, h and l, the unique equilibrium is for the seller to ask for a price that converges to l. If the price is not converging to l the same time β_t converges to 0, the buyer should wait for dt for a discrete drop of prices from P_T to l.

$\eta(t) = P_t - l$. We can then write for some non-zero constants c_1, c_2, c_3:

$$\int_0^T A(\tau)d\tau = c_1 \int_0^T \frac{1}{\eta(\tau)}d\tau \tag{61}$$

$$\int_0^T B(\tau)d\tau = \lim_{t\to T} \ln|\eta(t)| + c_2 \int_0^T \frac{1}{\eta(\tau)}d\tau + c_3 \tag{62}$$

Given that $\eta(t)$ converges to zero, we get a contradiction: either $\int_0^T A(\tau)d\tau$ or $\int_0^T B(\tau)d\tau$ is unbounded as required.

4.2 Equilibria with two-sided private information about fundamentals

Finally, to demonstrate that the result of delay is indeed a consequence of higher order beliefs and the possible exclusion of types, rather than simply the consequence of introducing two-sided private information, we briefly discuss equilibria with two-sided private information about fundamentals. Consider the following setup:

Seller \ Buyer	l	h
c_1	$1-\alpha, \gamma$	$\alpha, 1-\beta$
c_2	$1-\alpha, 1-\gamma$	α, β

(63)

where now the types of the seller denote the cost, with:

$$h > l > c_1 > c_2$$

here $l > c_1$ implies common knowledge of gains from trade. Note that we have assumed that the beliefs of the two types for the seller are identical in order to focus on the effects of two-sided private information about fundamentals without higher order uncertainty. Assume that

$$\alpha < \frac{l - c_2}{h - c_2}$$

This assumption guarantees that once type c_2 is revealed then she immediately offers $P_t = l$, hence the analysis is simplified.[16] In an equilibrium

[16]A similar result can be proven for beliefs that do not satisfy this condition, but it requires additional steps. The complication is the same as in the proof of Theorem 1 (compare the proof in the text with the proof in the appendix).

satisfying PO and RD, this lower cost seller type is randomizing between offering l and mimicking c_1 by offering P_t. Using the same arguments as the case of one-sided completely mixed beliefs, we can show that for a given δ she will reveal herself with probability 1 in finite time, $\bar{T}$. From that moment on, only the type c_1 is left and as there is a gap, the bargaining will end in a uniformly bounded (in δ) number of rounds, $\bar{\bar{n}}$. Furthermore, mirroring the reasoning in the proof of Theorem 2, the equilibrium price path $\{P_t\}$ has to be such that at $\bar{T}$ the price is at most $P_{\bar{T}} = h - \delta^{\bar{n}-1}(h-l)$ and the prices t before $\bar{T}$ have to satisfy:

$$P_{\bar{T}-t} < P_{\bar{T}} + t(h-l)\frac{1-\delta}{\delta}$$

Following the reasoning at the end of the proof of Theorem 2 the Coase conjecture holds for all such equilibria. We conclude that with common knowledge of gains from trade, in all the equilibria satisfying PO and RD the trade is efficient in the limit, as frictions disappear.

5 Final Remarks

In the previous section we have demonstrated that *possible exclusion of types* is a key element of the equilibria with delay constructed in Theorem 1. The intuition comes from the reputation literature. It seems that the main distinctive feature is the behavior of the different types once they are revealed. If all types have fully mixed beliefs, then once they are revealed they follow different price paths, but all those paths converge to the same limit as $\delta \to 1$, namely to offering l immediately. Therefore, for high discount factors there is little incentive to build reputation and delay disappears. In contrast, when one of the types is informed that the buyer has a high value, once revealed she can keep offering a high price. So even as $\delta \to 1$ the uninformed types behave differently from the informed one, hence incentives for building reputation arise. The same reasoning holds for the case of two-sided private information about fundamentals with common knowledge of strictly positive gains from trade: once the seller's type is revealed the gap leads her to offer $P_t = l$ (almost) immediately.

Does it imply that our result is restricted to models in which it is feasible for the seller to learn the buyer's value exactly? The answer is no. All that is necessary is *possible exclusion of types*, so that there is a positive probability that the seller has a type that puts probability zero on the lowest possible realization of the buyer's value. Stated differently, all that is required is that the buyer believes the seller *might* be informed about his valuation. Our results generalize further. Delay will occur even if the seller does not have a type that actually *knows* the valuation of the buyer – it suffices that she has a type that excludes (assigns zero probability) a low valuation. In a model in which the buyer has only two possible types these conditions coincide. But, for example, in a model in which the buyer has one of three possible values: $\{l, m, h\}$ (with $l < m < h$) for the existence of equilibria with non-trivial delay, it is sufficient that there is a positive probability that the seller has a type that puts probability 0 on the buyer being of type l. Using the same intuition, once revealed, a type with fully mixed beliefs offers price l (as $\delta \to 1$) while the type that puts probability 0 on l offers a price m. So even in the limit the two types behave differently once revealed, allowing for incentives to build reputation and non-trivial delay emerges.

It is also interesting to point out the similarities and contrasts between our equilibrium behavior and the one constructed in Abreu and Gul (2000). In their model the players can be behavioral/irrational with some probability (cf. Myerson, 1991) following the principle laid down by Kreps, Milgrom, Roberts and Wilson (1982). These irrational types follow a specific strategy. The rational players then are induced to follow mixed strategies imitating the irrational behavior and increasing the posterior of being irrational. In contrast, our framework assumes that *all* possible types are rational. Since the addition of the so-called irrational types is equivalent to the addition of rational types with different preferences, our framework can be seen as the addition of rational types with identical preferences but with a change in the information structure of the game. Schematically the work of Abreu and Gul (2000) can be seen as replicating a sub-tree of the original game tree where at the sub-tree a player is constrained to follow a specific strategy (the irrational behavior). The game begins by the selection of the original game or replicated sub-tree, and information sets are rearranged to reflect uncertainty as to the opponent type. We, on the other hand, replicate the whole game tree but rearrange the information sets in the replica (the possibly informed types). The game begins with a choice between the information structures and the information sets are rearranged according to the uncertainty about the uncertainties.

Finally, in our model we have assumed that the beliefs the buyer has over the possible beliefs of the seller are common knowledge. In other words, we have allowed only for second order uncertainties. This assumption is quite strong: if we point out that it is reasonable to think that the seller's beliefs are private, shouldn't the buyer's beliefs about the seller's beliefs be private as well? We certainly think that these higher order uncertainties are worthwhile pursuing. We treat our result as a first step in the more general analysis of higher order uncertainty in bargaining as well as in other models with asymmetric information.

6 Appendix

We provide the proof of delay in agreement for the general case where the uninformed seller's initial beliefs α_1 are arbitrary.

Proof of Theorem 1 for the case $\alpha_1 \geq l/h$. We construct the equilibrium in a similar way to the equilibrium constructed in the case $\alpha_1 < l/h$. The difference is that once the seller decides to deviate from the price P she will not immediately jump to an offer of l but will continue by offering a sequence of prices (each with probability one).

Following the rational of the proof of Theorem 1, we note that if the buyer type is revealed to be h at time t then the game ends with h paying P. Furthermore, if at time t the seller is revealed to be uninformed then the game continues according to the one-sided information one-sided offers bargaining game as described in Fudenberg and Tirole (1991, section 10.2.5 page 408).

The strategies in subgames in which it is common belief that the seller has type U are now more complicated than in the proof for $\alpha_1 < l/h$ and it makes the proof for the general case more complicated. Those strategies are uniquely determined and we now follow Fudenberg and Tirole (1991) to describe them: the seller follows a path of declining offers for $n(\alpha_t, \delta)$ periods with the buyer mixing and the seller eventually offering l after these $n(\alpha_t, \delta)$ periods if no previous price is accepted. The first (and maximal) price asked by the seller is given by $m = h - \delta^{n(\alpha_t,\delta)-1}(h-l)$. The Coase conjecture holds: for every α there exists a finite N such that $n(\alpha, \delta) < N$ for all δ. In particular, since α_t is not increasing, we can choose N such that $n(\alpha_1, \delta) < N$ for all δ which implies that whenever the seller reveals herself to be uninformed the price she asks for is no more than $h - \delta^N(h-l)$ and decreases to l within no more than N periods. We denote the seller expected payoff in this subgame by $F(\alpha_t, \delta)$. It is strictly increasing and continuous in

α_t.[17] It clearly satisfies:

$$(h - \delta^N(h-l)) \geq F(\alpha_t, \delta) \geq l$$

Finally it can be shown that $\lim_{\delta \to 1} \frac{\partial F(\alpha_t,\delta)}{\partial \alpha_t} = 0$.[18]

If the seller is revealed to be type I^h he offers price $p_t = h$ and it is immediately accepted. That finishes the description of strategies in subgames in which the seller's type is revealed.

Now, suppose that we are in a stage during which there is still uncertainty about the type of the seller. Consider a price P strictly between l and h. Type I^h will be choosing P at every time t. Type I^l will be choosing l at every period. Type l will be accepting only offers up to (including) l. Seller U will be mixing between asking P and revealing herself with another offer. Buyer h will be mixing between accepting P and rejecting it.

If the seller is mixing at time t she must be indifferent between mimicking I^h and revealing herself. Her payoff today if she offers P is $\alpha_t\mu_t P + \delta(1 - \alpha_t\mu_t)F(\alpha_{t+1}, \delta)$ where $F(\alpha_{t+1}, \delta)$ is her expected payoff tomorrow, given that the buyer rejected P at time t and hence the seller's updated beliefs are given by α_{t+1} (and next period she will be mixing again). If the seller is mixing today we must have:

$$\alpha_t\mu_t P + \delta(1 - \alpha_t\mu_t)F(\alpha_{t+1}, \delta) = F(\alpha_t, \delta) \tag{64}$$

rewriting we get

$$\alpha_t\mu_t = \frac{F(\alpha_t, \delta) - \delta F(\alpha_{t+1}, \delta)}{P - \delta F(\alpha_{t+1}, \delta)} \tag{65}$$

using (2) and (65) we obtain:

$$\alpha_t = \left(1 - \frac{F(\alpha_t, \delta) - \delta F(\alpha_{t+1}, \delta)}{P - \delta F(\alpha_{t+1}, \delta)}\right)\alpha_{t+1} + \frac{F(\alpha_t, \delta) - \delta F(\alpha_{t+1}, \delta)}{P - \delta F(\alpha_{t+1}, \delta)} \tag{66}$$

[17] While the seller's expected payoff is continuous in α_t, the buyer's expected payoff is not, but in the construction of equilibria we only use continuity of $F(\alpha_t, \delta)$.

[18] If we increase α_t by a small $\Delta\alpha$, the most the seller can gain is if he increases the price to p_{t-1} (instead of p_t) and all this additional mass accepts it. That gives him at most a gain of $\left(p_{t-1} - \delta^N l\right)\Delta\alpha$. So the derivative is bounded by $\left(p_{t-1} - \delta^N l\right) \to 0$.

Consider a lower bound on probabilities that the seller assigns to type h for which she will be willing to choose P, i.e. assume that h will accept P with probability one and find beliefs so that the seller would offer P :

$$\alpha_t P + \delta(1-\alpha_t)F(\alpha_{t+1},\delta) \geq F(\alpha_t,\delta)$$

but α_{t+1} will be equal to 0 (since the buyer accepts with probability one) and $F(0,\delta) = l$. So the bound is a solution to $\alpha^* P + \delta(1-\alpha^*)l = F(\alpha^*,\delta)$ i.e.,

$$\alpha^* = \frac{F(\alpha^*,\delta) - \delta l}{P - \delta l} \tag{67}$$

The RHS is continuous in α^*. For $\alpha^* = 0$ the RHS is positive. For $\alpha^* < 1$ the limit of the RHS as $\delta \to 1$ is 0 (because $F(\alpha^*,\delta)$ converges to l). So for large δ equation (67) has a solution. Let α^* be the maximal solution to (67).

There are a few things important to notice about equation (66). First, for large enough δ so that $F(\alpha_1,\delta) < P$, for any $\alpha^* < \alpha_t \leq \alpha_1 < 1$ the solution to this equation exists and satisfies $\alpha_{t+1} < \alpha_t$, as the equation states that α_t is a convex combination of 1 and α_{t+1}. Second, for large δ the RHS is strictly increasing in α_{t+1}, so there is a unique solution for each α_t. Finally, (again for large δ) the solution is strictly increasing in α_t.

Similarly, if the buyer is randomizing, he is indifferent between accepting the offer P at time t or rejecting and getting the expected payoff at time $t+1$. The payoff at time $t+1$ is P if the seller offers P, or the payoff is $h - m = \delta^{n(\alpha_{t+1},\delta)-1}(h-l)$ if the seller reveals herself to be uninformed at time $t+1$. Hence we have

$$h - P = \delta((1-\beta_{t+1}\sigma_{t+1})(h-P) + \beta_{t+1}\sigma_{t+1}\delta^{n(\alpha_{t+1},\delta)-1}(h-l)) \tag{68}$$

$$\beta_t\sigma_t = \frac{(1-\delta)(h-P)}{\delta^{n(\alpha_t,\delta)-1}(h-l) - \delta(h-P)} \tag{69}$$

using Bayes rule (3) and condition (69) we have:

$$\beta_t = \left(1 - \frac{(1-\delta)(h-P)}{\delta^{n(\alpha_t,\delta)-1}(h-l) - \delta(h-P)}\right)\beta_{t+1} + \frac{(1-\delta)(h-P)}{\delta^{n(\alpha_t,\delta)-1}(h-l) - \delta(h-P)} \tag{70}$$

Let

$$\beta^* = \frac{(1-\delta)(h-P)}{\delta^{n(\alpha^*,\delta)-1}(h-l) - \delta(h-P)} \tag{71}$$

If at time t we have $\alpha_t = \alpha^*$, for every $\beta < \beta^*$ the buyer will be better off accepting P (even if next period the U seller reveals herself), and for every $\beta > \beta^*$ there is a probability $\sigma_t < 1$ such that the buyer wouldn't mind waiting another period.

Now we can construct the equilibrium strategies: Start with finding α^*. Set $\alpha(1) = \alpha^*$ and using (66) find $\alpha(2)$ as the largest solution α_t to this equation (with $\alpha_{t+1} = \alpha(1)$). Continue in this fashion to find a sequence $\alpha(n)$. For δ large enough so that $F(\alpha_1, \delta) < P$ we get an increasing sequence: $\alpha(n) < \alpha(n+1)$. Given this sequence we can use (70) to construct a similar sequence $\beta(n)$ with $\beta(1) = \beta^*$. Now, if the prior beliefs (α_1, β_1) are on the path $(\alpha(n), \beta(n))$ then we have found an equilibrium: the mixing probabilities can be found using (65) and (69).

If they are not on the path, then we describe the strategies in the first round as follows (analogously to the proof in the text for the simple case). Define $N(\alpha_1) = \max\{n|\alpha(n) < \alpha_1\}$ and $M(\beta_1) = \max\{n|\beta(n) < \beta_1\}$. Since α_1, β_1 are given we denote $N = N(\alpha_1)$ and $M = M(\beta_1)$.

There are two possible cases:

1. $N > M$
2. $N \leq M$

In the first case (the easier one) we define the following strategies. Let $\sigma_1 = 0$ and σ_t satisfy (69) for all $2 \leq t \leq M+1$. Let μ_1 be such that $\alpha_2 = \alpha(M)$ as derived from (2). Since

$$\alpha_1 \geq \alpha(N) > \alpha(M) \tag{72}$$

(note the strict inequality), the U seller is strictly better off by mimicking the I^h seller in the first round, so $\sigma_1 = 0$ is a strict best response. For $2 \leq t \leq M$ let μ_t follow (65) and α_t follow $\alpha(n)$. As $\beta_1 = \beta_2$ and the sequence of α_t is defined, we can use (70) to find the sequence of β_t and (69) to find the sequence of σ_t. At period $t = M$ the buyer is still mixing since $\beta_M > \beta^*$ and this leads to the seller beliefs being $\alpha_{M+1} = \alpha^*$. Since $\beta_{M+1} < \beta^*$, the buyer will accept P at time $M+1$ which is exactly why a seller at $\alpha_{M+1} = \alpha^*$ will be indifferent and can follow $\sigma_{M+1} = 1$ as required.

In the second case ($N \leq M$) we construct the following strategies. At time 1, the seller will be mixing with high probability so that the posterior *after* that offer will be on a path leading to β^*. The buyer will be mixing starting at α_1. The difficulty is that after we use (66) to find the path of α_t this path is different than $\alpha(n)$ so that we have to find a new sequence of

$\beta(n)$. Also, the definition of β^* has to change to depend on α_{N+1} instead of α^*.

Formally, given $\alpha_1' = \alpha_1 \in (\alpha(N), \alpha(N+1))$ equation (66) defines a sequence α_t' with $\alpha_t' \in (\alpha(N-t+1), \alpha(N-t+2))$. At $t = N$, $\alpha_N' > \alpha^*$ and at $t = N+1$, $\alpha_{N+1}' < \alpha^*$. Comparing with the sequence $\alpha_t = \alpha(N-t+2)$ we have $\alpha_t' < \alpha_t$. Now define $\beta^{'*}$ to satisfy:

$$\beta'^* = \frac{(1-\delta)(h-P)}{\delta^{n(\alpha_{N+1}',\delta)-1}(h-l) - \delta(h-P)}$$

as the critical belief that makes the buyer indifferent between accepting and rejecting P even if he knows that next period $\alpha_{N+1}' < \alpha^*$ so that the seller if uninformed will reveal herself for sure. Given that $\alpha_{N+1}' < \alpha^*$ we get $\beta'^* \leq \beta^*$. Then, using (70)we define a sequence $\beta'(n)$ using α_t' instead of α_t, Unfortunately given that the weights are now different (with smaller weight on β_{t+1}) we cannot guarantee that $\beta'(n) \leq \beta(n)$. Therefore we also cannot guarantee that $M' = \max\{n | \beta'(n) < \beta_1\}$ is larger than N. If it is, then we have found an equilibrium. If it is not larger (and we conjecture that this will never be a case for large enough δ), then we do not know how to construct the equilibrium.

Finally, the bargaining can continue with positive probability for up to $Min\{M(\delta), N(\delta)\}$ periods. Denote by $\underline{w}_\alpha$ and $\overline{w}_\alpha$ the smallest and the largest weights in (66) put on 1. Then (approximately)

$$\frac{\ln(1-\alpha_1)}{\ln(1-\underline{w}_\alpha)} \leq N(\delta) \leq \frac{\ln(1-\alpha_1)}{\ln(1-\overline{w}_\alpha)}$$

$$\begin{aligned} \underline{w}_\alpha &\geq \left(\frac{(1-\delta)\,l}{P-\delta l}\right) \\ \overline{w}_\alpha &\leq \frac{F(\alpha_1,\delta)-\delta l}{P-\delta l} \end{aligned}$$

so that for large δ both $\frac{1}{\ln(1-\underline{w}_\alpha)}$ and $\frac{1}{\ln(1-\overline{w}_\alpha)}$ are of the order of $\frac{1}{\ln\delta}$. So that N grows at the same rate as the frequency of offers.

Similarly, define $\underline{w}_\beta$ and $\overline{w}_\beta$ and note that:

$$\underline{w}_\beta \geq \frac{(1-\delta)(h-P)}{\delta^{n(\alpha_1,\delta)-1}(h-l)-\delta(h-P)}$$

$$\overline{w}_\beta \leq \frac{(1-\delta)(h-P)}{(h-l)-\delta(h-P)}$$

to obtain the same conclusion for M. ■

References

[1] Abreu D. and F. Gul (2000). "Bargaining and Reputation," Econometrica 68(1): 85-117.

[2] Admati A. R. and M. Perry (1987). "Strategic delay in bargaining," Review of Economic Studies 54: 345–364.

[3] Ausubel L. and R. Deneckere (1998). "Reputation in bargaining and durable goods monopoly," Econometrica 57, 511-531.

[4] Cho I. K. (1990). "Uncertainty and delay in bargaining," Review of Economic Studies 57, 575-596.

[5] Cramton P. C. (1984). "Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty," Review of Economic Studies 51: 579-593.

[6] Cramton P. C. (1992). "Strategic delay in bargaining with two-sided uncertainty," Review of Economic Studies 59: 205-225.

[7] Deneckere R. and M. Liang (2001). "Bargaining with interdependent values," Working paper.

[8] Evans R. (1989). "Sequential bargaining with correlated values," Review of Economic Studies 56: 499–510.

[9] Kreps D., P. Milgrom, J. Roberts, and R. Wilson (1982). "Rational cooperation in the finitely repeated prisoners' dilemma," Journal of Economic Theory 27: 245-252.

[10] Kreps D. M., and R. Wilson (1982). "Reputation and Imperfect Information," Journal of Economic Theory 27: 253-279.

[11] Fudenberg D., D. Levine and J. Tirole (1985). Infinite-horizon models of bargaining with one-sided incomplete information. In Game-Theoretic Models of Bargaining, ed. A. Roth. Cambridge University Press.

[12] Fudenberg D. and J. Tirole (1991). Game Theory. Cambridge, MA. MIT Press.

[13] Gul F., H. Sonnenschein and R. Wilson (1986). "Foundations of Dynamic Monopoly and the Coase Conjecture," Journal of Economic Theory 39: 155-190.

[14] Milgrom P., and J. Roberts (1982) "Predation, Reputation and Entry Deterrence," Journal of Economic Theory 27: 280-312.

[15] Myerson R. (1991). Game Theory: Analysis of Conflict. Cambridge, MA. Harvard University Press.

[16] Vincent D. R. (1989) "Bargaining with common values," Journal of Economic Theory 48: 47–62.

[17] Yildiz M. (2001). "Sequential Bargaining without a Common Prior on the Recognition Process," working paper, MIT.