ELSEVIER

Economic Design 1 (1994) 79–102

Economic Design

On the possibility of efficient bilateral trade

Sanjeev Goyal [1]

Econometric Institute, Erasmus University, 3000 DR Rotterdam, Netherlands

Abstract

I study a sequential process in which different pairs of traders bargain over the terms of trade of an indivisible good. I consider both *one-sided* and *two-sided* offers based bargaining at the stage-game level. The sequential process is modelled as an infinite stage-game of incomplete information and the paper studies the efficiency properties of its equilibria. It is shown: With one-sided offers, all equilibria are long-run ex post efficient; with two-sided offers, examples of equilibria are constructed with widely varying efficiency properties.

Key words: Bilateral bargaining; Learning; Ex post efficiency

JEL Classification: C78, D82, D83

1. Introduction

In several economic transactions – examples include housing contracts, used car sales, labour agreements – the terms of trade are determined through bilateral bargaining between individual agents who are (mutually) imperfectly informed about their opponent's valuation for the good. An interesting question in such

[1] Econometric Institute, Erasmus University, Post Box 1738, 3000 DR Rotterdam, The Netherlands. Telephone 31-10-452-4082269; Fax: 31-10-452-7746. This paper is based on Chapter 1 of my Ph.D thesis. I would like to thank my thesis committee members, Larry Blume, Tapan Mitra, Karl Shell, and specially, David Easley, for encouragement and advice. The paper has benefitted greatly from the suggestions of the editor and an associate editor of this journal and from the detailed comments of an anonymous referee. Thanks are also due to Venkatesh Bala, Steven Brahms, Birgit Grodal, Patrick Legros, Eric Maskin, Frans Meershoek, Pauline Rutsaert, Larry Samuelson, Jonathan Thomas and John Wooders for helpful comments.

SSDI 0928-5040(94)00005-T

contexts is whether the bilateral bargaining is efficient i.e., if it realizes all the potential gains from trade. For bargaining between any *single* pair of agents this question has been answered in the negative by Myerson and Satterthwaite (1983). They showed that with two-sided incomplete information there does not exist any trading rule / mechanism which can realize ex post efficient outcomes.

In this paper, I locate the bilateral bargaining between any two agents within a sequence of similar bargaining interactions between different traders. This sequentiality of trading creates the possibility that agents can access historical data concerning bargaining between similar traders in the past. A good example of this possibility is the bargaining between a buyer and a seller in the housing market; typically, both sides check if similar houses have been recently bought and sold and if so, at what price, before proceeding with their own bargaining. I study the role of such information transmission (and the consequent learning) in the process of bilateral bargaining. More specifically, I ask if some standard one-sided offers and two-sided offers trading mechanisms realize all gains from trade in the long run, as a consequence of learning from the past?

The sequence of bargaining interactions is formally studied as an infinite stage-game of incomplete information. Bargaining rules based on *one-sided* offers as well as *two-sided* offers are considered, at the stage game level. The one-sided offers game is modelled via the take-it-or-leave-it offer and the two-sided offers bargaining is based on simultaneous sealed bids. The principal results concern learning and ex post efficiency of Bayes-Nash equilibria in the infinite game. An equilibrium is said to be long-run (LR) ex post efficient if for all stages after some finite stage, outcomes correspond to the ex post efficient outcomes. A related issue of interest is the nature of long-run learning of true valuations. I shall say that complete learning occurs when traders' beliefs assign (roughly speaking) point mass on the true valuation of their opponent. In what follows, I use the term generic to denote events that happen with probability 1 given a trader's prior belief concerning his opponent's valuation.

In the game with one-sided offers it is shown that every equilibrium is LR ex post efficient (Theorem 2) and that learning by the active side of the market is incomplete, generically (Theorem 3). In the one-sided offers case, the focus is on learning by the side of the market that makes the offers. Valuations of traders are said to be compatible if the buyer's valuation is strictly greater than the valuation of the seller; and valuations are said to be incompatible, if the buyer's valuation is strictly less than that of the seller. Proposition 2.1 and Lemma 2.2.1 show that when valuations are incompatible learning is always incomplete. Proposition 2.2 shows that if valuations are compatible then, after a finite stage, trade takes place at the same price forever. This implies that in such cases learning *ceases* after a finite stage. Since, at any finite stage, the support of a trader's posterior belief has a non-degenerate support, this also shows that in any equilibrium, for such valuations, learning will be incomplete. These results put together yield Theorem 3: in any equilibrium learning is incomplete, generically.

For the game with two-sided offers bargaining the analysis is done via a set of examples. These examples illustrate the range of possible outcomes that can obtain in the equilibria of the infinite stage game. I begin by constructing equilibria with little or no learning, and poor efficiency properties. On closer scrutiny it becomes clear that the inefficiency in these equilibria arises out of the 'stand-off' nature of the stage game strategies: traders make bids/offers which they know, given their posterior beliefs, will result in trade with zero probability. This observation helps me to construct an equilibrium with the LR ex post efficiency property. It is also shown that in this equilibrium, learning is generically incomplete. These examples suggest that in the two-sided offers case, though LR ex post efficiency[4] is attainable, one cannot hope for the strong general results obtained in the one-sided offers case.

The results of this paper are related to several strands of the literature. I now discuss these relationships briefly.

In recent years the study of the welfare properties of decentralized trading has been an active area of research, (see, for example, Symposium on Non-Cooperative Bargaining in The Journal of Economic Theory (1989), the survey paper by Wilson (1987) and the papers cited therein). An important component of this research has been the search for efficient mechanisms. A mechanism defines, in this setting, a game of incomplete information, which is usually solved using the concept of Bayes-Nash Equilibrium. For a mechanism to be efficient one requirement might be that all Bayes-Nash equilibria of the game thus defined are (in a suitable sense) efficient. An even stronger requirement might be that such efficiency obtains for a *class* of environments (cf. uniform efficiency, Wilson (1987)). I explore both these requirements in this paper. In particular, Theorem 2 establishes that in markets with sequential trading, the take-it-or-leave-it offer trading mechanism fulfils both these requirements. Taken together with the examples for the two-sided offers case this suggests that the one-sided offers based bargaining rule has very attractive welfare properties.

In this connection, it is worth mentioning a recent paper by Ausubel and Deneckere (1993), who show (roughly speaking) that for the single good bargaining problem the one-sided offers trading rule can approximately attain the ex-ante efficient expected utilities. The analysis in my paper concerns welfare properties of different trading rules in the context of an infinite sequence of one-shot bargains. In this context, the main result for one-sided offers bargaining, Theorem 2, shows that the take-it-or-leave-it offer bargaining has good welfare properties. This result can be seen as strengthening the case for such one-sided offers trading rules, as well as providing some basis for their widespread use.

The results on incomplete learning are of independent interest. In particular, Theorem 3 is closely related to results in Aghion et al. (1991). Aghion et al. (1991) consider an example with a long-lived (dynamic optimizing) monopoly selling repeatedly to a myopic consumer whose valuation is private information and distributed uniformly over the unit interval. They show that learning will typically be incomplete, and with positive probability it will stop at some finite point. I examine a model in which both sides of the market are myopic; however, the distribution of the buyer's valuation is no longer restricted to be uniform and is quite general. Theorem 3 shows that if valuations are compatible then learning *ceases* at some finite point and long-run learning will, therefore, be incomplete, with probability 1. In view of the existing results on incomplete learning, this is not a particularly surprising result; the incentives to experiment and hence learn are lower for myopic agents as compared to long-sighted agents. The results reported here, however, help clarify the relationship between efficiency and learning in deterministic multi-agent settings. The existing literature suggests that incomplete learning, though *ex ante* optimal, may be 'inadequate' or undesirable since it might prevent the attainment of ex post optimal outcomes. In contrast, Theorem 3 and Example 3 show that for a general class of markets with bilateral trading, 'aggregate' *ex post* efficiency obtains, though learning is generically incomplete. The paper may also be interpreted as a study of the long-run behaviour of a sequence of one-shot Bayes-Nash equilibria. An interesting issue is: Does this sequence converge to Nash play with respect to the 'true game', i.e., the stage-game with the actual valuations of the traders? There is a large literature on the subject of learning to play equilibrium, (see the surveys by Blume and Easley (1992,1993) and the papers referred to therein). In particular, Jordan (1991) examines this issue for finite action games in which players know their own pay-offs but do not know their opponent's pay-offs. He shows that under some conditions this convergence does occur. The examples in this paper suggest that in the context of infinite action stage-games, convergence of strategies to Nash equilibrium strategies may not reflect any learning of pay-offs of the true game. In particular, Example 1 in the discussion on two-sided offers is revealing: there is no learning and yet the (stationary) Bayes-Nash strategies constitute a Nash equilibrium with respect to any true game. On the other hand, Example 3 presents an equilibrium in which there is learning, and the sequence of one-stage Bayes-Nash equilibria also converges to some Nash equilibrium of the true game, for all possible such games.

The paper is organized in five parts. In Section 2 the model is formally developed. Section 3 presents the general results with one-sided offers and section 4 presents three examples for two-sided offers based bargaining. Section 5 concludes. The proofs are given in the appendix at the end of the paper.

2. The model

Consider pairs of traders, comprising of one buyer and one seller, meeting to trade an indivisible good at discrete points in time. Both traders have personal valuations for the good which are private knowledge. These valuations are identical across the stages for all buyers and all sellers, respectively. Traders know this fact, and have priors on these valuations. Thus in every stage-game a new pair of traders enters, with prior beliefs about opponent's valuations, observes outcomes of past bargaining, updates these prior beliefs and then bargains over the terms of trade. This bargaining is based on either one-sided offers or two-sided simultaneous offers. One-sided offer bargaining takes the form of the *take-it-or-leave-it* offer game: one side makes an offer, if the other side accepts the offer then they exchange the good, otherwise not. In either case that ends the transaction. In the two-sided offers case traders *simultaneously submit a sealed-bid.* The good is exchanged if the buyer's bid exceeds that of the seller's, at a price midway between the two bids; if the bids are incompatible no trade is realized. The traders then exit. The outcome of this bargaining is observed by the next stage traders and the process begins afresh.

This structure is somewhat special – in particular, if different traders enter each stage then the assumption that they all have identical valuations seems to be quite restrictive. The analysis of such a model is nevertheless interesting because, I believe, the insights that one can obtain in this simple setting also carry over to more realistic settings. As an illustration consider the following setting: In each period a pair of different traders meet to bargain over the terms of trade of an indivisible good. The valuations of the sellers are drawn in each stage from some finite set of valuations whereas the (privately known) valuations of the buyers are identical over time. Suppose that traders are perfectly informed about the valuations (of their predecessors) on their side of the market. Also assume that the sellers side of the market makes the offers and the buyers respond by accepting or rejecting the offers. In such a setting, a seller learns from the history of trades not just of sellers with the same valuation as himself but also from those with different valuations. The learning and efficiency properties of the equilibria of the game defined by such a trading process can, however, be analyzed using the same methods as those reported in this paper. I now briefly outline how this is done.

At any stage, the learning from history is reflected in the posterior distribution of the seller. The seller makes an offer which optimizes one-period pay-offs given these beliefs and his own valuation. Using arguments in section 3.3 it is possible to show that, eventually, sellers with valuations lower than the buyers' valuation will settle down on one offer which is accepted by buyers. The seller types with valuations higher than the buyers' valuation, on the other hand, will keep making offers higher than the buyers' valuation and no trade will occur. The principal results on LR ex post efficiency and generic incomplete learning, Theorems 2 and 3, thus extend, in a natural way to this setting.

I next briefly discuss the case in which the valuations of the buyers side (i.e., the 'responding side' of the market) are allowed to vary, whereas those on the sellers side are identical, over time. There are different ways in which this may be formulated. One possibility is the following formulation: [2] There is a myopic

[2] I am grateful to the associate editor for suggesting this formulation.

seller making offers to short lived buyers with valuations drawn from some distribution. The interesting situation here is one in which the sellers do not know the distribution of the buyers' valuations (note that if the distribution is known to the sellers then there is little that they can learn by observing past trades, since the valuations are randomly drawn in every period). The problem of learning the parameters of an unknown distribution is, however, a special case of the single agent learning problem which has been extensively studied in the literature. A well known result in this area concerns incomplete learning of the parameters of the distribution by a long lived agent (for the relevant references see the surveys by Blume and Easley (1992), (1993)). In this context, it is worth noting that since the objective of the seller is the maximization of the one-period expected profits, even if he knew the true distribution, the optimal offer will typically be in the interior of the support of the buyers' distribution. This suggests that, in this setting, the appropriate notion of ex post efficiency is different from the one we consider in the present paper (cf. Definition 2.2. below).

An alternative formulation of varying buyers' valuations is the following: The buyers start off with a privately known valuation which changes at some point to some new privately known valuation. Assume that the myopic seller knows the general structure of the process but does not know the initial valuation, the point of change or the final valuation. In such a setting, in addition to the questions concerning learning and ex post efficiency there is also the interesting issue of how the historical experience affects the behaviour of the offers sequence (even after the change in buyers' valuations). In some simple examples (and for some specific valuations) I have been able to show that ex post efficiency as well as incomplete learning obtain. A general treatment of this model, however, falls outside the scope of the present paper.

In passing, I should add that an alternative interpretation of the model in its present form is that it is an infinitely repeated game with two-sided incomplete information between two myopic players.

I now proceed with a formal description of the trading process as an infinite stage-game of incomplete information (Harsanyi, 1967-8).

2.1. The game

1. Order of Play: Denote stages by $n = 1,2,3,\dots$. At each stage a different pair of traders, comprising of one buyer and one seller, meets to trade an indivisible good. I refer to the bargaining between the traders at the n^{th} stage, as the n^{th} *stage game.*

2. Two-sided Incomplete Information: Traders have personal valuations for the good: v_b for the buyer and v_s for the seller. These valuations are respectively identical across the stages. Buyers (sellers) consider v_s (v_b) to be a random variable. v_b and v_s are assumed to be stochastically independent. Each buyer (seller) on entry has a prior distribution F (G) over the v_s (v_b). As in Myerson and Satterthwaite (1983), I assume that F and G are continuously differentiable, the density functions exist, are continuous and that $F'(x) = f(x) > 0$, and $G'(x) = g(x) > 0$, for all $x \in [0,1]$. The minimal assumptions on F and G suggest that the class of environments under consideration is fairly general.

3. Bargaining Procedure: Two different bargaining procedures are considered: One-sided offers and two-sided offers. In the one-sided offer case a trader makes a take-it-or-leave-it offer to which the opponent responds with an acceptance or a rejection. On acceptance, the good is exchanged at the price defined by the offer; on rejection, no exchange takes place. Irrespective of the outcome, traders exit. In the analysis below we speak, for expositional simplicity, in terms of the seller's offer game and refer to it as the *seller's game*; the entire analysis can be reproduced in terms of the buyer's offer game. Denote by s_n the offer by the seller at the n^{th} stage. The buyer's decision is either accept(a) or reject(r). On acceptance of s_n the good is exchanged at price of $p_n = s_n$ and the pay-offs to the buyer and seller are $(v_b - p_n)$ and $(p_n - v_s)$, respectively.

In the two-sided offer case, traders submit offers b_n and s_n. The good is exchanged if $b_n \geqslant s_n$, at a price, $p_n = (b_n + s_n)/2$. The pay-offs are as defined in the one-sided offer case. We refer to this procedure as the *simultaneous sealed bid game.*

The formal analysis of the one-sided offers game (in section 3) considers the case in which the same side of the market makes the offers in all the stages. I believe that the arguments presented in that case can also be used to analyze the case where, across the stages, different sides of the market make take-it-or-leave-it offers. [3]

4. Preferences: Traders are risk-neutral and maximize expected earnings. Denote by π_n^s and π_n^b the expected earnings for the seller and the buyer, respectively. These pay-offs are determined by b_n and s_n with the expectations conditioned on the *posterior* beliefs of the traders, F_n and G_n, respectively. To understand the evolution of these posterior beliefs, we next describe the flow of information across stages.

5. History: For any stage n, the observed set is $H = \{T, NT\}$, where $T = \{trade\}$ and $NT = \{no\ trade\}$. Denote by h_n the actual outcome; thus $h_n = \{T\}$ or $\{NT\}$,

[3] To see how this can be done, consider the example with uniform priors presented in section 3.1. Assume, for simplicity, that the sellers (buyers) side makes offers in odd (even) numbered stages. Also assume that traders observe the offers and the trade outcomes of all previous stages. I now sketch the argument for the efficiency result in this setting. From the discussion in section 3.1 it follows that the sellers side completely reveals its valuation through the first period offer. Thus in subsequent stages, buyers will always be able to trade by making an offer equal to the valuation of the sellers; the stages in which the sellers make an offer can be analyzed along the lines of section 3. These observations put together yield the efficiency result.

and define history at stage n, as $h^n = \{h_k\}_{k=1}^{n-1}$ where $h^n \in H^{n-1} = H \times H \times \ldots \times H$. It is assumed that stage 1 traders start from zero history. (More generally, from now on, subscripts denote single period outcomes and superscripts the sequence until that stage. I also use $h^n = \{NT\}$ *or* $\{T\}$, to denotes histories where $\forall k < n$, $h_k = \{NT\}$ or $\{T\}$, respectively.) Information conveyed across time is thus rather minimal and one can imagine a large class of environments/markets which will satisfy this requirement.

Richer information transmission can be modelled easily, within this framework. A natural candidate for such a more informative history would be one which incorporates the sequence of actual prices, p_n, at which trade occurs. The results on efficiency and learning presented below extend to such histories in a natural way.

6. Strategies: Seller's game: At any stage n, a seller's *pure* strategy is a function, s_n, that maps his personal valuation, v_s, and the history of past outcomes, h^n, into the set of the possible offers that he can make, given by [0,1]. Formally, s_n: $[0,1] \times H^{n-1} \to [0,1]$. The buyer's pure strategy, b_n, is a function that maps his personal valuation, v_b, and the offer he receives from the seller, s_n into the set of possible actions for him, given by $\{a,r\}$. Formally, b_n: $[0,1] \times [0,1] \to \{a,r\}$.

Simultaneous sealed bid game: In the two sided offers case both the buyer and the seller make offers and the buyer's pure strategy is defined exactly as the seller's strategy, i.e., b_n: $[0,1] \times H^{n-1} \to [0,1]$. I restrict myself to pure strategies in this paper.

For stage n, denote by S_n^b and S_n^s the pure strategy set of the buyer and the seller, respectively. Also denote by $S_n = S_n^b \times S_n^s$ the strategy set for the n^{th} pair. The strategy set until stage n is then defined as $S^{n-1} = \prod_{k=1}^{n-1} S_k$. Denote the strategy pair for traders at stage game n by $t_n = (b_n,\ s_n)$.

7. Beliefs: At each stage a buyer and a seller enter with private valuations, v_b and v_s, respectively, and prior belief about their opponent's valuations, denoted by F and G, respectively. They know (in equilibrium) the strategies of previous traders. After entry, they observe the history of past play, h_n. Given his private valuation, the buyer (seller) can then reconstruct the offers made by his predecessors, and the possible responses of the corresponding sellers (buyers). Using this knowledge of offers, the resulting outcomes and hence the set of potential responses, he then updates his priors with the help of Bayes' Rule.

Let $d_n = \{F_n,\ G_n\}$; the following notation concerning the support of the belief distributions is extensively used: $u_g^n(v_s)$ denotes the supremum and $l_g^n(v_s)$ refers to the infimum, respectively of the support of G_n. $u_f^n(v_b)$ and $l_f^n(v_b)$ denote the corresponding values for the beliefs of the buyer. The support of the limiting beliefs, $(G^*,\ F^*)$, is defined in terms of $u_g^*(v_s)$ and $l_g^*(v_s)$ for the seller and $u_f^*(v_b)$ and $l_f^*(v_b)$ for the buyer. To avoid excessive notation, posterior beliefs are not shown to depend on private valuations; I will on occasion also omit mention of these valuations while discussing the support of beliefs.

Bayes' Rule cannot be applied when a player encounters zero probability events, e.g., histories which could have happened along the equilibrium path, only with zero probability. I assume that when faced with such histories traders use a pre-specified rule, which is discussed below (see Definition 2.1).

2.1.1. The equilibrium concept

An equilibrium is an infinite sequence of strategy-belief pairs, $\{t_n, d_n\}_{n=1}^{\infty}$, such that for every stage n, given any history, h^n, the beliefs of the traders, d_n, are well defined, and the strategy pair for the stage game, t_n, constitutes a Bayes-Nash equilibrium with respect to these beliefs.

Updated beliefs are derived using Bayes' Rule as long as a player observes histories that occur with positive probability along the equilibrium path. I next consider the evolution of beliefs off the equilibrium path. In my model, due to the two-sided nature of incomplete information, an off-the-equilibrium-path history may not be immediately detected as a deviation, by both the players. This problem is further complicated by the minimal information that is transmitted across stages. In view of the absence of any strategic link across stages this issue is, however, not central to the analysis. I therefore adopt a simple convention when players learn that the previous history is off the equilibrium path: when it is common-knowledge that a history is inconsistent with equilibrium strategies then both players choose to ignore any information the history contains and revert to their prior beliefs. All traders know this rule and so subsequent players can use the truncated history starting at this stage game.

Formally, I denote histories that are consistent with equilibrium strategies as consistent histories, and those that are inconsistent with equilibrium strategies as inconsistent histories. If it is common knowledge at stage n that a history is inconsistent then players, at subsequent stages, use the truncated history from stage n onwards. I refer to such histories as n stage onward consistent histories. When a history is privately known to be inconsistent but this inconsistency is not common knowledge, then no restriction is imposed on the beliefs.

It may seem as though the complications in this formalization arise due to the assumption that traders do not observe prices at which previous trades take place. To some extent this is true; for instance, in the one-sided (seller's) offer case, the observation of price offers would immediately reveal the deviations by the seller. It is easy to see, however, that the deviations/mistakes made by a buyer are not immediately revealed even if the price offers are observed. Thus observing the sequence of price offers helps but does not resolve all the difficulties. (For a discussion of the complications that arise during the updating process in the two-sided offers case, see section 4 below).

The discussion so far is summarized in the following definition.

Definition 2.1. An equilibrium is a sequence of strategy-belief pairs, $\{\hat{t}_n, \hat{d}_n\}_{n=1}^{\infty}$, such that:

1. Given any particular $\hat{d}_n$, the strategy pair for the n^{th} stage game, $\{\hat{s}_n, \hat{b}_n\}$ satisfy,

$$(i)\,\pi_n^b(\hat{b}_n(v_b, h^n), \hat{s}_n \mid \hat{F}_n) \geqslant \pi_n^b(b, \hat{s}_n \mid \hat{F}_n), \forall b \in [0,1], \forall v_b \in [0,1].$$

$$(ii)\,\pi_n^s(\hat{s}_n(v_s, h^n), \hat{b}_n \mid \hat{G}_n) \geqslant \pi_n^s(s, \hat{b}_n \mid \hat{G}_n), \forall b \in [0,1], \forall v_s \in [0,1],$$

2. The beliefs of the traders, $\hat{d}_n = (F_n, G_n)$, are derived using $\{\hat{t}^n, h^n\}$, as follows: (a) If it is common knowledge that h^n is *inconsistent* w.r.t. $\hat{t}^n$, then $F_n = F$, $G_n = G$, $\forall v_b, v_s \in [0,1]$; If it is privately but not commonly known that a history is inconsistent then no restriction is imposed on beliefs of a player.

(b) If h^n is $k < (n-1)$ onward consistent w.r.t., $\hat{t}^n$ then, for a seller by an application of Bayes' Rule,

$$G_n(x) = \frac{G(x) - G(l_g^n)}{G(u_g^n) - G(l_g^n)}, \text{ for } x \in [l_g^n, u_g^n];$$

$$G_n(x) = 1, \text{ for } x > u_g^n, \text{ and } G_n(x) = 0, \text{ for } x < l_g^n$$

A similar procedure defines F_n.

It is worth discussing two properties of the updating formula presented above. One, implicit in the formula for the updated posterior belief is the assumption that equilibrium strategies are (weakly) monotonic with respect to personal valuations – this generates the interval support for the posterior belief. It is clear from Lemma 1.1 below that for the one-sided offers case this is not a restriction. For the two-sided offers case the examples considered in section 4 all satisfy the monotonicity property and so the formula suffices for the purposes of this paper. Two, I assume in 2(a) that faced with zero probability events, traders revert to original priors. Alternative specifications can easily be considered. One such alternative would be to have traders revert to the posterior belief which they held immediately before the discovery of the zero probability event. This formulation retains more of the already accumulated information than the one presented in the paper. A moment's thought, however, reveals that such a modification in the specification has no significant effect on the main results.

2.1.2. Long-run (LR) ex post efficiency

Given the results of Myerson and Satterthwaite (1983), one cannot hope to find an equilibrium that is ex-post efficient (with any bargaining rule) for all stages of the infinite stage game. This motivates the use of a weaker efficiency requirement: long-run ex post efficiency.

For the single bilateral trading problem, ex post efficiency requires that a good is transferred from a seller to a buyer if $v_b > v_s$, and not traded if $v_b < v_s$. If $v_b > v_s$ and $h = \{NT\}$, then loss of welfare is $(v_b - v_s)$, likewise if $v_b < v_s$, and $h = \{T\}$, then welfare loss is $(v_s - v_b)$. Note that in my model, in both bargaining cases, at any stage either all or none of the gains from trade are realized. This allows me to define a long run version of ex post efficiency:

Definition 2.2. An equilibrium is said to be LR ex post efficient if, for every compatible pair (v_b, v_s), there is some $N(v_b, v_s)$, s.t., $h_k = \{T\}$ $\forall k \geqslant N$, and if for every incompatible pair (v_b, v_s), there is some finite $N'(v_b, v_s)$, s.t., $h_k = \{NT\}$, $\forall k \geqslant N'(v_b, v_s)$.

Remark: This definition is much stronger than long-run average ex post efficiency.

3. One-sided offers, learning and efficiency

The main results is this section, Theorems 2 and 3, may be summarized as follows: If bargaining proceeds on the basis of one-sided offers then every equilibrium of the infinite game is LR ex post efficient and exhibits incomplete learning, generically. Theorem 1 establishes the existence of an equilibrium.

In the seller's game if $v_b > v_s$ one expects that, in the long-run, sellers would extract all the potential surplus, i.e., the limit of the sequence of prices, $\lim_{n \to \infty} p_n = p^* = v_b$. Proposition 2.2 establishes the somewhat surprising result that, after some finite point, trade takes place at the same price forever, and this price is, typically, strictly lower than the true valuation of the buyer. This incomplete surplus extraction by the seller is one aspect of a more general phenomenon: long-run learning by the seller, of the buyer's valuation, is generically incomplete. I begin with an example which illustrates the intuition behind these general results.

3.1. An example with uniform priors

This section illustrates the general results of the next two sections within the context of an example where both F and G are uniform distributions. Recall, in the one-sided offers bargaining attention is restricted to learning by the active side of the market.

To fix ideas, assume that actual realizations of the valuations are: $v_b = 15/32$; $v_s = 1/4$. Given the requirement of sequential rationality (see Definition 3.1. below), it is easy to show that for stage 1, the unique first stage equilibrium is given by: $s_1(v_s = 1/4) = 5/8$, and $b_1(v_b = 15/32,\ s_1 = 5/8) = \mathrm{r}$. Thus, $h^2 = \{NT\}$. The beliefs of the seller at stage 2 may be computed using Bayes' Rule; and the optimal offer of the second stage seller is given by $s_2(v_s = 1/4,\ h^2 = \{NT\}) = 7/16$. Given that $v_b = 15/32$, the buyer's optimal response, $b_2(v_b = 15/32,\ s_2 = 7/16) = a$. Thus $h_2 = \{T\}$ and $h^3 = \{h_1 = NT;\ h_2 = T\}$.

I next show that in this setting, once trade occurs, it occurs forever after. The reason is rather simple. The unique optimal offer, for the seller with valuation $v_s = 1/4$, given the history h^3, is $s_3(v_s = 1/4,\ h^3) = 7/16$. Given that b_n is stationary, it follows that $h_3 = \{T\}$.

Since no additional information will be revealed by the outcome in the bargaining at the 3^{rd} stage game, the unique optimal offer for stage 4 will be $s_4(v_s = 1/4,\ h^4) = s_3(v_s = 1/4,\ h^3)$, and $b_4(v_b = 15/32,\ s_4 = 7/16) = a, \ldots$ and so on. Similar reasoning applies to all such pairs. Moreover, given the description of equilibrium behaviour, no trade takes place if $v_b < v_s$. Thus, the equilibrium is LR ex post efficient.

This example also gives us an idea of why learning will typically be incomplete. In particular, since optimal offers remain stationary after $n = 3$, and trade always occurs learning also ceases; in other words, $u_g^* = 5/8$ and $l_g^* = 7/16$. Theorem 3 is a generalization of this insight.

It should be mentioned here that, in general, following on a realized trade, the optimal offer is not unique, and this creates the possibility of $h_{k+1} = \{NT\}$, despite $h_k = \{T\}$, (even along the equilibrium path!). The main result on efficiency for one sided offers, Theorem 2 below, allows for multiple optimal continuation offers, and also takes into account off the equilibrium path play.

3.2. *Existence of equilibrium*

Given the sequential move structure of the *take-it-or-leave-it offer game*, it is natural to refine the Nash requirement at the stage-game to allow for *sequential rationality*. I require that given the beliefs of the seller about the buyer's valuations, G_n, the strategy pair, $\{\hat{s}_n,\ \hat{b}_n\}$, constitutes a Nash equilibrium of the n^{th} stage game that satisfies sequential rationality for the buyer. Also note that given the specification of the seller's game, observation of history is relevant for the offers of the seller only, a fact that greatly simplifies the analysis. In particular, we require that the seller simply starts afresh, whenever he faces what he regards as a privately inconsistent history. Definition 2.1 can now be reformulated as follows:

Definition 3.1. An equilibrium is a sequence of strategy-belief pairs, $\{(\hat{t}_n\ \hat{G}_n)\}_{n=1}^{\infty}$, such that for any n, the following is true.
1. Given $\hat{G}_n$, the strategy pair for the n^{th} stage game, $\{\hat{s}_n,\ \hat{b}_n\}$ satisfies,

$$(i)\,\pi_n^b\big(\hat{b}_n(v_b,\ s)\big) \geqslant \pi_n^b(b,\ v_b,\ s),\ \forall b \in [0,1],\ \forall s \in [0,1],\ \forall v_b \in [0,1]$$

$$(ii)\,\pi_n^s\big(\hat{s}_n(v_s,\ h^n),\ \hat{b}_n \mid \hat{G}_n\big) \geqslant \pi_n^s\big(s,\ \hat{b}_n \mid \hat{G}_n\big),\ \forall s \in [0,1],\ \forall v_s \in [0,1]$$

2. The beliefs of the seller, $\hat{G}_n$, are derived from $\{\hat{t}^n,\ h^n\}$, as follows: (a) If the seller knows that h^n is *inconsistent* w.r.t. $\hat{t}^n$ then, $\hat{G}^n = G$.

(b) If h^n is $k < (n-1)$ onward consistent w.r.t. $\hat{t}^n$, then by an application of Bayes' Rule,

$$\hat{G}_n(x) = \frac{G(x) - G(l_g^n)}{G(u_g^n) - G(l_g^n)} \text{ for } x \in [l_g^n, u_g^n] \text{ and}$$

$$\hat{G}_n(x) = 1, \text{ for } x > u_g^n, \text{ and } \hat{G}_n = 0, \text{ for } x > l_g^n$$

The idea of sequential rationality (within the stage-game) is incorporated in the requirement that the buyer's response is optimal with respect to all $s \in [0,1]$, not just the equilibrium offer s_n. It will be assumed, in the discussion of the seller's game that when faced with an offer $s_n = v_b$, the buyer always accepts.

The first step in the analysis is an existence theorem. The proof of existence involves showing that there exists an infinite sequence of strategy-belief pairs which satisfies, for every finite stage n, certain conditions. To this end, the first step is a characterization of the sequentially rational strategies for the n^{th} stage game, given any beliefs G_n. It is shown that equilibrium strategies for the stage game have a particularly simple structure: a buyer at any stage accepts an offer if and only if $v_b \geqslant s_n$. (Lemma 1.1. in the appendix). Next I consider a sequence $\{t_n, G_n\}_{n=1}^{\infty}$, where the beliefs satisfy condition 2 in Definition 3.1. To demonstrate existence I have to show that for every n, a strategy pair, t_n, satisfying condition 1, in Definition 3.1, exists.

Theorem 1. There exists a sequence of strategy-belief pairs, $\{t_n, G_n\}_{n=1}^{\infty}$, such that for any n, Conditions 1 and 2 in the equilibrium definition are satisfied.

3.3. Learning and efficiency

An equilibrium is LR ex post efficient if trade occurs for all stages after some finite stage when $v_b > v_s$, while trade does not occur after some finite stage if $v_b < v_s$. I begin by establishing that any equilibrium has the latter property.

Proposition 2.1. If $v_b < v_s$ then in any equilibrium continuation, $h_k = \{NT\}$, $\forall k \geqslant n$.

The proof consists of showing that $s_n(v_s) \geqslant v_s$ for all v_s, at all stages in an equilibrium continuation. Given the description of buyer's strategy this implies that if $v_b < v_s$ then $h_k = \{NT\}$, for all stages.

I next consider the situation where the valuations are compatible, i.e., $v_b > v_s$. I show that in any equilibrium continuation, $h_k = \{T\}$, after some finite stage. Proposition 2.2 establishes the claim for continuations from inconsistent and zero histories and Proposition 2.3 considers all other histories.

There are two steps in the proof of Proposition 2.2. First, it is shown that if $v_b > v_s$ then there is some finite $N(v_b, v_s)$, such that $h_N = \{T\}$; second, it is shown that there is some $\hat{N} \geqslant N(v_b, v_s)$, such that $\forall k \geqslant \hat{N}$, $h_k = \{T\}$. (In the uniform priors case, we saw that once an offer has been accepted then at all subsequent stages it is the unique optimal offer for the seller. This is not true in general, and that motivates the second step in the analysis.)

The intuition behind the first step can be seen in the example with uniform priors: For every stage in which trade does not occur, at the following stage the seller truncates the support of the distribution of the buyer's valuation. If trade never occurs then this process of truncation has the following general implication:

Lemma 2.2.1. $u_g^*(v_s, h^\infty = \{NT\}) = v_s$ and $l_g^*(v_s, h^\infty = \{NT\}) = 0$.

Given $h^\infty = \{NT\}$, it follows that for every $k \geqslant n$, $s_k(\cdot) = u_g^{k+1}(\cdot)$. However, by assumption, $v_b > v_s$, and so it follows from Lemma 2.2.1 that there must exist some stage $k^* \geqslant n$ in which the optimal offer of the seller $\hat{s}_{k^*}(v_s, h^{k^*} = \{NT\}) \leqslant v_b$. Given the requirement of sequential rationality in Definition 3.1., this implies that the buyer accepts the offer for this stage and that contradicts the hypothesis that trade never occurs. Thus if $v_b > v_s$ then there is some finite stage k such that $h_k = \{T\}$.

The argument in the second step builds on this result and shows that if $v_b > v_s$ then the outcome 'no trade' cannot occur infinitely often. If no trade occurs then the process of truncation is active, and leads over time to a situation in which the probability mass gets concentrated on a smaller range of values of the buyer's valuation. In particular, if no trade occurs infinitely often (henceforth denoted by i.o.) then, in the limit, I obtain the following:

Lemma 2.2.2. Suppose $v_b > v_s$. Then $u_g^*(v_s, h_k = \{NT\}\ i.o.) = l_g^*(v_s, h_k = \{NT\}\ i.o) = v_b$.

(Lemmas 2.2.1 and 2.2.2 appear to be inconsistent. To see that they are not, note that in the latter it is assumed that $v_b > v_s$ and this implies (from the above arguments) that there is some k such that $h_k = \{T\}$; in other words, if $v_b > v_s$ then in equilibrium $h^\infty = \{NT\}$ is ruled out.)

The essential trade-off in the optimization problem of the seller is between the higher payoff that results from a higher bid (if it is successful) as against the lower probability of its acceptance. Lemma 2.2.2 shows that if 'no trade' occurs very often then the probability mass will get concentrated over a smaller and smaller support; this implies that over time the trade-off moves increasingly in favour of a bid closer to the lower bound of the support of the posterior beliefs of the sellers. The proof of Proposition 2.2 formalizes this intuition and shows that at some finite point, the loss from the lowering of probability of trade becomes sufficiently high to make the infimum of the support, $l_g^k(v_s)$, the optimal bid of the seller. In equilibrium play since $v_b \geqslant l_g^k(v_s)$, this results in *trade* and also implies that there is no truncation of support, at the subsequent stage. Hence, the optimal bid remains the same from that point onwards, so does the seller's bid and also the trade outcome. I can now state,

Proposition 2.2. Suppose $v_b > v_s$. In any equilibrium continuation, from a history h^n s.t., $G_n = G$, there is some $N(v_b, v_s) \geqslant n$, such that for all $k \geqslant N(v_b, v_s)$, $h_k = \{T\}$.

Proposition 2.3 extends the above result to the case of some stage onward consistent histories.

Proposition 2.3. Suppose $v_b > v_s$. In any equilibrium continuation from history h^n which is k (where $k < n$) onward consistent, there is some $N(v_b, v_s) \geqslant n$ such that $\forall k \geqslant N(v_b, v_s)$, $h_k = \{T\}$.

(The arguments in the proof of Proposition 2.3 are similar to those presented in the proof of Proposition 2.2. and are omitted; for details see Goyal (1991)).

The above results are obtained in the setting of the seller's game; since the buyer's game is symmetric, they are also true for the corresponding buyer's game. This allows me to state the following general result.

Theorem 2. All equilibria with one-sided offers based bargaining are LR ex post efficient.

The interest now turns to the nature of long-run learning in such equilibria. For the case where $v_b < v_s$, Proposition 2.1 and Lemma 2.2.1 together imply that learning is always incomplete. On the other hand, if $v_b > v_s$, arguments from Proposition 2.2 show that after some finite stage, since trade takes place at the same price no further learning is possible. Given that at any finite stage, the support of a seller's belief distribution is non-degenerate, this allows me to state the following result.

Theorem 3. In any equilibrium with one-sided offers based bargaining, learning is incomplete, generically.

It is worth adding that this incomplete learning is not due to the assumption that sellers cannot observe prices at which previous trades have taken place. To see why this is the case it is instructive to look at the example presented in section 3.1: If the seller observes the price at which trade takes place at $n = 2$, it would make no difference to his posterior belief at $n = 3$ and hence no difference to the optimal offer by him at $n = 3$.

4. Two-sided offers, learning and efficiency

The analysis of the two-sided offers bargaining model is done in terms of examples. The idea of these examples is to illustrate the wide range of possibilities

– from extreme inefficiency to LR ex post efficiency – in different equilibria of the infinite game when bargaining is based on two-sided offers.

Example 1: (Extreme inefficiency)

$$\forall n, \{s_n = 1, \forall v_s \in [0,1];\ b_n = 0, \forall v_b \in [0,1]\}.$$

It is easy to show that this sequence of strategy pairs can be sustained in an equilibrium. In such an equilibrium, no trade or learning ever takes place.

Example 2: (Some learning, some gains from trade)

$$\forall n, \{b_n = 0, \forall v_b \in [0,1/2); = 1/2, \forall v_b \in [1/2,1]\}.$$

$$\forall n, \{s_n = 1/2, \forall v_s \in [0,1/2]; = 1, \forall v_s \in (1/2,1]\}.$$

It is possible to show that this strategy sequence can be sustained in an equilibrium. In this equilibrium, at stage 1, if $v_s \leqslant 1/2 \leqslant v_b$, then $h_1 = \{T\}$; otherwise not. Is there any learning in this equilibrium? If $v_b \geqslant 1/2$, and $h_1 = \{T\}$, then $u_f^2(v_b) = 1/2$ and $l_f^2(v_b) = 0$. On the other hand if $v_b > v_s > 1/2$, and $h_1 = \{NT\}$, then from the knowledge of equilibrium strategies (and personal valuations), the buyer at stage 2 does learn that $v_s > 1/2$. However, this learning is not usefully employed, by him or any later buyer since they all bid $b_n = 1/2$, despite knowing that the probability of trade is zero.

It is this 'unreasonableness' of these buyer bids that motivates the construction of the following equilibrium, which is LR ex post efficient. Details of the proof showing that this is an equilibrium of the infinite stage game are given in the appendix.

Example 3: (LR ex post efficient equilibrium)

Consider the following strategies-beliefs sequence:

Strategies:

For stage 1,

$$\hat{b}_1(v_b) = 0, \text{ for } v_b \in \left[0, \frac{1}{2}\right); = \frac{1}{2}, \text{ for } v_b \in \left[\frac{1}{2}, 1\right].$$

$$\hat{s}_1(v_s) = \frac{1}{2}, \text{ for } v_b \in \left[0, \frac{1}{2}\right]; = 1, \text{ for } v_b \in \left(\frac{1}{2}, 1\right].$$

For stage n,

If $h^n = \{NT\}$,

$$\hat{b}_n(v_b, \cdot) = 0, \text{ for } v_b \in \left[0, \frac{1}{2^n}\right); = \frac{1}{2^n}, \text{ for } v_b \in \left[\frac{1}{2^n}, \frac{2}{2^n}\right); \ldots;$$

$$= \frac{2^n - 1}{2^n}, \text{ for } v_b \in \left[\frac{2^n - 1}{2^n}, 1\right].$$

$$\hat{s}_n(v_s, \cdot) = \frac{1}{2^n}, \text{ for } v_s \in \left[0, \frac{1}{2^n}\right]; = \frac{2}{2^n}, \text{ for } v_s \in \left(\frac{1}{2^n}, \frac{2}{2^n}\right]; \ldots;$$
$$= 1, \text{ for } v_s \in \left(\frac{2^n - 1}{2^n}, 1\right].$$

If h^n is s.t. for some $1 \leqslant k \leqslant (n-2)$, $h_k = \{NT\}$, and $h^{n/k} = \{T\}$, then $\hat{b}_n(v_b, h^n) = \hat{b}_{n-1}(v_b, h^{n-1})$, and $\hat{s}_n(v_s, h^n) = \hat{s}_{n-1}(v_s, h^{n-1})$.
If h^n is s.t. $h_{n-2} = \{T\}$, but $h_{n-1} = \{NT\}$, then $\hat{b}_n(v_b, h^n) = \hat{b}_1(v_b)$, $\hat{s}_n(v_s, h^n) = \hat{s}_1(v_s)$.
If h^n is commonly known to be $k \leqslant n-1$ onward consistent, then $(\hat{b}_n, \hat{s}_n)$ are defined using the rules above w.r.t. histories $h^{n/k}$.
Beliefs: For stage 1, $F_1 = F$, $G_1 = G$; For stage n, if h^n is consistent then F_n and G_n are derived via Bayes' Rule, using all history. For stage n, if h^n is s.t. $h_{n-2} = \{T\}$, but $h_{n-1} = \{NT\}$, then history is commonly known to be inconsistent and $F_n = F$ and $G_n = G$. For h_n which is $k \leqslant n-1$ onward consistent, F_n and G_n are derived via Bayes' Rule, using k onward history $h^{n/k}$. Finally, if h^n is s.t. $h_{n-1} = \{T\}$, and some type of some player knows, (but this is not commonly known), that $\text{Prob}_{F_{n-1}}(h_{n-1} = \{T\}) = 0$, or $\text{Prob}_{G_{n-1}}(h_{n-1} = \{T\}) = 0$, then for that type, $F_n = F_{n-1}$, or $G_n = G_{n-1}$, depending on whether he is a buyer or a seller.

To see that this equilibrium is LR ex post efficient, let us consider a particular pair of realizations. Suppose that $v_b = 4/9$, $v_s = 3/9$. Then, $\hat{b}_1(v_b = 4/9) = 0$, and $\hat{s}_1(v_s = 3/9) = 1/2$, hence $h_1 = \{NT\}$. In stage 2, $\hat{b}_2(v_b = 4/9, h_1 = \{NT\}) = 1/4$ and $\hat{s}_2(v_s = 3/9, h_1 = \{NT\}) = 1/2$, hence $h_2 = \{NT\}$. In stage 3, $\hat{b}_3(v_b = 4/9, h^3 = \{NT\}) = 3/8$ and $\hat{s}_3(v_s = 3/9, h^3 = \{NT\} = 3/8$. The bids are compatible and $h_3 = \{T\}$. Given that $h^2 = \{NT\}$, and $h_3 = \{T\}$, $\hat{b}_4(\cdot)$ $\hat{b}_3(\cdot)$ and $\hat{s}_4(\cdot) = \hat{s}_3(\cdot)$, which implies $h_4 = \{T\}$, and so on....for all $k \geqslant 3$. Similar arguments work for other compatible pairs. Moreover, from the description of the equilibrium strategies it is clear that if $v_b < v_s$ then no trade occurs.

A closer scrutiny of the general structure of the stage-game equilibrium strategies reveals an interesting feature: If $v_b > v_s$ then in every stage at least one of the traders bids with some intention to trade. Formally, there is some $0 < \epsilon < 1/2$, such that

$$s_n(v_s) \leqslant u_g^n(v_s) - \epsilon\big(u_g^n(v_s) - v_s\big) \text{ and/or } b_n \geqslant l_f^n(v_b) + \epsilon\big(v_b - l_f^n(v_b)\big).$$

Moreover, in the above equilibrium, once trade occurs, traders persist with the same bids and offers. From the above example, it is easily seen that this has the implication of limiting learning to the stages before trade first takes place. In particular, for the valuations considered, for the sellers, $u_g^k = 1/2$ and $l_g^k = 3/8$ and, for the buyers, $u_f^k = 3/8$ and $l_f^k = 1/4$, for all $k \geqslant 3$. This argument is quite general and holds for all compatible pairs of valuations. If valuations are incompatible learning is naturally incomplete since no seller (buyer) makes a bid/offer lower (higher) than his own valuation. These observations put together suggest that learning in the equilibrium given in Example 3 will typically be incomplete.

Example 3 can be generalized to show that any equilibrium of the infinite game in which the stage-game strategies have these properties is LR ex post efficient and also exhibits incomplete learning, generically. The formal statements and proofs may be found in Goyal (1993).

Example 3 also helps us to illustrate a problem with the common knowledge of inconsistent histories (alluded to in section 2 above). Suppose that $v_s < v_b < 1/2$, and that at stage 1, $h_1 = \{T\}$. Given common knowledge of equilibrium strategies, at stage 2, the buyer knows that history is inconsistent, since $b_1(v_b < 1/2) = 0 < s_n(v_s)$, $\forall v_s \in [0,1]$. However, such a buyer also knows that a seller with $v_s \leqslant 1/2$ does not know that the history is inconsistent, since $b_1 \geqslant 1/2 = s_1(v_s \leqslant 1/2)$ is possible, given $\hat{b}_1$ and $G_1 = G$. In other words, inconsistency of history is not commonly known. This problem is intrinsic to the nature of two-sided incomplete information, and can arise even if bids/offers of previous traders were common knowledge.

5. Conclusion

Typically, in bilateral bargaining traders are incompletely informed about the true valuations of their opponent. Such incompleteness of information, it is well known, can result in bargaining outcomes which are inefficient. This paper shows that, in a general class of environments where trading is sequential, learning from historical data can help overcome this problem and lead to long run ex post efficient outcomes. The extent of long run learning about opponents' valuations is also examined. It is shown to be incomplete, generically.

Appendix: proofs

Fact I. Consider the sequence, $\{t_n, G_n\}_{n=1}^{\infty}$. Assume that $\{G_n\}$ satisfy condition 2, in the equilibrium definition, and t_k exists and is well defined $\forall k \leqslant n$. Then at stage $(n+1)$, $\forall h^{n+1} \in H^n$, support $G_{n+1} = [l_g^{n+1}, u_g^{n+1}]$, for some $0 \leqslant l_g^{n+1} \leqslant u_g^{n+1} \leqslant 1$.

Proof. The proof is in two parts: First it is shown that, for $v_s \in [0,1]$, $\forall k \leqslant n$, support $G_k = [l_g^k, u_g^k]$. Then it is shown that *support* $G_{n+1} = [l_g^{n+1}, u_g^{n+1}]$. An application of the principle of induction then completes the proof. The proof is straightforward and omitted. For details the reader may refer to Goyal (1991).

Lemma 1.1. Given beliefs, G_n, $\{s_n, b_n\}$ are sequentially rational strategies for the n^{th} stage game, if, and only if, they satisfy the following:

$$(i)\ s_n(v_s, h_n/G_n) = \operatorname{argmax}_s\left\{\int_s^{u^n}(s - v_s)\,dG_n(b),\ \forall s \in [l^n, u^n];\right.$$
$$\left.(s - v_s),\ \forall s < l^n; 0,\ \forall s > u^n\right\}$$

$$(ii)\ b_n(v_b, s_n) = a \text{ if } v_b \geqslant s_n; = r \text{ if } v_b < s_n.$$

Proof. The proof is straightforward and omitted.

Theorem 1. I use an argument based on induction. Clearly, b_n satisfying condition (ii) in lemma 1.1 always exists. The focus is on the nature of the pay-off function of the seller, $\pi_n^s(s_n(v_s, h^n), b_n/G_n)$. I show that, for any n, for each realization of the valuation of the seller, it is a continuous function, with respect to $s \in [0,1]$. An application of Weierstrass's Theorem, then completes the proof.

Stage Game 1: For any $v_s \in [0,1]$, $\pi_1^s((v_s, s), b_n \mid G) = (s - v_s)(G(1) - G(s))$, for $s \in [0,1]$. Given that G is continuously differentiable, $\pi_1^s(\cdot)$ is a continuous function for all $s \in [0,1]$. Hence, by Weirstrass's Theorem, $s_1(v_s)$ is well defined for all realizations of v_s.

Stage Game 2: The seller in the 2^{nd} stage game receives the history, h^2, from the first stage. In particular, $h^2 = \{T\}$ *or* $\{NT\}$. Suppose, that $h^2 = T$, then the *support* $G_n = [s_1(v_s), 1]$, by Bayes' updating. The pay-off to the seller can now be defined as follows,

$$\pi_n^s(s, b_2/G_2)$$
$$= \{(s - v_s)(G_2(1) - G_2(s)),\ \forall s \geqslant s_1(v_s);\ (s - v_s),\ \forall s < s_1(v_s)\} \quad (1)$$

Simplifying, one can see that,

$$(s - v_s)(G_2(1) - G_2(s)) = \frac{((s - v_s)(1 - G(s))}{(1 - G(s_1(v_s))} \quad (2)$$

Now, given the assumptions on G, it follows that π_2^s is a continuous function. Analogous arguments can be used to show continuity in the case that $h^1 = \{NT\}$. An application of Weierstrass Theorem then establishes that the maximand is well defined and obtained at some $s \in [0,1]$.

Stage n: Suppose that $\forall v_s \in [0,1]$, s_k is well defined for $k \leqslant n$. Consider the $n + 1^{th}$ stage.

Stage $(n + 1)$: Given that G_n satisfies condition 2 in Def 2.1. and s_k is well defined for $k \leqslant n$, from Fact 1 *support* $G_{n+1} = [l_g^{n+1}, u_g^{n+1}]$. If $[l_g^n, u_g^n] = [0,1]$,

s_{n+1} is well defined using arguments form stage 1. If not, i.e., h^{n-1} is consistent, and $0 < l_g^{n+1} < u_g^{n+1} < 1$, then, for $s \in [l_g^{n+1}, u_g^{n+1}]$,

$$\begin{aligned}
&\pi_{n+1}^s(v_s, s, b_{n+1}/G_{n+1}) \\
&\quad = \Big\{(s - v_s)\big(G_{n+1}(u_g^{n+1}) - G_{n+1}(s)\big), \forall s \in \big[l_g^{n+1}, u_g^{n+1}\big]; \\
&\qquad (s - v_s), \forall s < l^{n+1}; 0, \forall s > u^{n+1}\Big\} \qquad (3)
\end{aligned}$$

Simplifying,

$$(s - v_s)\big(G_{n+1}(u_g^{n+1}) - G_{n+1}(s)\big) = (s - v_s)\frac{\big(G(u_g^{n+1}) - G(s)\big)}{G(u_g^{n+1}) - G(l_g^{n+1})} \qquad (4)$$

It is easy to check now that $\forall v_s \in [0,1]$, π_{n+1}, is a continuous function w.r.t. $s \in [0,1]$, $\forall v_s \in [0,1]$. An application of Weierstrass' Theorem then completes the proof. *Q.E.D.*

Proposition 2.1. Recall that $\forall k \geqslant n$, $b_k(v_b, s_k) = b(v_b, s_k) = a$, $\forall v_b \geqslant s_k; = r$, $\forall_b < s_k$. To establish the result then I only have to show that seller's optimal strategy always prescribes, $s_k(v_s, \cdot) \geqslant v_s$, for $k \geqslant n$. When $u_g^n(v_s) \geqslant v_s \geqslant l_g^n(v_s)$, this is immediate from the rationality of the seller; given the nature of history, however, it cannot be otherwise. That completes the proof. *Q.E.D.*

Lemma 2.2.1: Fix some realization of the seller's valuation, v_s. First, it is shown that u_g^* exists. Given the hypotheses, (old history) h^n can be ignored by traders starting at stage n. Since $\forall k > n$, traders make equilibrium bids and offers, Bayes' updating implies that $u_g^{k+1} \leqslant u_g^k$. Since $v_b \in [0,1]$, the sequence $u_g^k(v_s)$ is monotonic and bounded between 0 and 1. Hence a limit to the sequence exists and $u_g^*(v_s)$ is well defined.

Next it is shown that $u_g^*(v_s) = v_s$. The proof is by contradiction. Suppose not. From Proposition 2.1, since $u^k(v_s) \geqslant v_s$, $\forall k$, it follows that $u_g^*(v_s) \geqslant v_s + \delta$, for some $\delta > 0$. By hypothesis, $h_k = \{NT\}$, $\forall k \geqslant n$. Hence, given that $G_n = G$, and this is an equilibrium continuation, $u_g^k(v_s) = s_{k-1}(v_s)$, $\forall \geqslant n$ and thus for *any* $\epsilon > 0$, there is some $N(\epsilon)$, such that $\forall k \geqslant N(\epsilon)$, $u_g^*(v_s) \leqslant s_k(v_s) \leqslant u_g^*(v_s) + \epsilon$. It is shown that this contradicts the optimality of s_k for some $k \leqslant N(\epsilon)$.

Take a point x, where $v_s < x < u_g^*(v_s)$ such that $G(u^*) - G(x) = \hat{g} > 0$. The pay-off to the seller at any stage, $k \geqslant n$ is given by

$$\pi_n^s(v_s, x, h^k = \{NT\}) = (x - v_s)\big(G_k(u_g^{k-1}) - G_k(x)\big) \geqslant (x - v_s)\,\hat{g} > 0 \qquad (5)$$

On the other hand, for any $k \geqslant n$, the pay-off to an equilibrium offer, $\hat{s}_n(v_s)$, will be given by

$$\pi_k^s\left(s_k\left(v_s,\ h^k=\{NT\}\right)\right)=\left(s(v_s)-v_s\right)\left(G_k\left(u_g^k\right)-G\left(s_k(v_s)\right)\right)$$
$$\leqslant(1-v_s)\frac{G\left(u_g^k\right)-G\left(u_g^*\right)}{G\left(u_g^*\right)} \tag{6}$$

Since G is continuous, $\lim_{k\to\infty} G(u_g^k)=G(u_g^*)$. Hence, $\lim_{k\to\infty} \pi_k^s(s_k(v_s,\ h^k=\{NT\})\to 0$. The pay-off from offering x is bounded away from 0, and so this contradicts the optimality of $\hat{s}_n(v_s)$ for some $k>N(\epsilon)$. Thus, $u_g^*(v_s,\ h^\infty=\{NT\}) = \hat{v}_s$. *Q.E.D.*

Lemma 2.2.2. First, note that the existence of limits follows from standard arguments, as elaborated in Lemma 2.2.1 above. Since this is a sequence of beliefs generated by equilibrium play, starting from $G_n=G$, Bayes' updating implies that $l_g^k(v_s)\leqslant v_b$, and $u_g^k(v_s)\geqslant v_b$ for $k\geqslant n$. Thus $l_g^*(v_s)$ and $u_g^*(v_s)$ exist and $l_g^*(v_s) \leqslant v_b \leqslant u_g^*(v_s)$.

I next show that when $h_k=\{NT\}$, i.o., $u_g^*(v_s)=l_g^*(v_s)=v_b$. Suppose not, then given that $u_g^*\geqslant v_b$ and $l_g^*\leqslant v_b$, it must be the case that $u_g^*(v_s)\geqslant l_g^*(v_s)+\delta$, for some $\delta>0$. By hypothesis, $h_k=\{NT\}$, i.o. for $k\geqslant n$. Since $u_g^k(v_s)\to u_g^*(v_s)$, for any small $\epsilon>0$, there is some $k(\epsilon)$, such that $\forall k\geqslant k(\epsilon)$, $u_g^k(v_s)\leqslant u_g^*(v_s)+\epsilon$. Now we can, using arguments very similar to Lemma 2.2.1, show that there is some bid lying strictly between $l_g^*(v_s)$ and $u_g^*(v_s)$ which will, *for all stages* after some stage, k, strictly dominate any bid $s_k(v_s)>u_g^*(v_s)$. If $h_k=\{NT\}$ i.o., then for some such stage despite $s_k<u_g^*(v_s)$, $h_k=\{NT\}$, implying that $u_g^{k+1}(v_s)< u_g^*(v_s)$. It follows from Bayes' formula that $u_g^n(v_s)$ is (at least weakly) monotonically decreasing across time, this contradicts the hypothesis that $u_g^k(v_s)\to u_g^*(v_s)$. *Q.E.D.*

Proposition 2.2. Given that *support* $G_k=[l_g^k,\ u_g^k]$, (From Fact 1, in the appendix), for a seller making an offer $s_k(v_s)\in[l_g^k(v_s),\ u_g^k(v_s)]$, the pay-off is:

$$\pi_n^s\left(s_k\left(v_s,\ h^k\right),\ b_k/G_k\right)=\left(s_k(v_s)-v_s\right)\frac{G\left(u_g^k\right)-G\left(s_k(v_s)\right)}{G\left(u_g^k\right)-G\left(l_g^k\right)} \tag{7}$$

Differentiating and simplifying, w.r.t. the offer of the seller s, we get,

$$\pi_n'^s(v_s,\ s,\ b_k/G_k)\leqslant 1-\frac{g(s)\cdot(s-v_s)}{G\left(u_g^k\right)-G\left(l_g^k\right)},\ \forall s\in\left(l_g^k(v_s),\ u_g^k(v_s)\right) \tag{8}$$

From Lemma 2.2.2 above, $lim_{k\to\infty}u_g^k(v_s)=v_b=lim_{k\to\infty}l_g^k(v_s)$, and so it follows from the continuity of G and g that there is some k^*, such that $\forall s_n\in[l_g^{k*},\ u_g^{k*}]$, $\pi_n'^s((v_s,\ s),\ b_n\,|\,G_{k^*})<0$. Thus the optimal offer for the seller $s_{k^*}(v_s)=$

$l_g^{k^*}(v_s)$. Given that $l_g^{k^*}(v_s)$ is equal to an accepted offer at some stage $k < k^*$ in the past, this implies that at stage k^* in equilibrium, $\hat{b}_{k^*}(v_b, \hat{s}_{k^*}(v_s) = l_g^{k^*}) = a$. Bayes' updating then implies that $G_{k^*+1} = G_{k^*}$. Hence, at stage $k^* + 1$ too, $\hat{s}_{k^*+1}(v_s) = l_g^{k^*}(v_s)$ and $h_{k^*+1} = \{T\}$. An application of the principle of induction completes the proof. *Q.E.D.*

Example 3. (Proof that the strategies and beliefs are an equilibrium.) First, we consider the beliefs. So long as history is fully (or some stage onward) consistent, the sequence of beliefs, $\{F_n, G_n\}$, is derived by Bayes' Rule, thus satisfying requirement (2)(b) in Definition 2.1. If history is commonly known to be inconsistent, then (F_n, G_n), satisfy the rule for such histories, stated in Definition 2.1. If history is privately but not commonly known to be inconsistent then, since no restriction applies, the above construction meets the requirement of Definition 2.1.

The strategies are now examined. First note that given some history, monotonicity of strategies w.r.t., valuations follows from description of equilibrium.

Stage 1: Consider $\hat{b}_1$. Suppose $v_b \in [0,1/2)$, then given $F_1 = F$ and $\hat{s}_1$, $\pi_1^b(\hat{b}_1(v_b) = 0,\ \hat{s}_1/F) = 0$. For any bid $b \in [0,1]$, $\pi_1^b(1/2 > b > 0,\ \hat{s}_1/F) = 0$, and $\pi_1^b(b \geqslant 1/2,\ \hat{s}_1/\mathrm{F}) < 0$. This establishes optimality of $\hat{b}_1$ for $v_b \in [0,1/2)$. For $v_b \in [1/2,1]$, $\pi_1^b(\hat{b}_1(v_b) = 1/2,\ \hat{s}/F) = (v_b - 1/2)\mathrm{Prob}_F(v_s \leqslant 1/2) > 0$; For other possible bids $b \in [0,1]$, $\pi_1^b(b_1(v_b) < 1/2,\ \hat{s}_1/F) = 0$; $\pi_1^b(b > 1/2,\ \hat{s}_1/F) = (v_b - (1/2 + b)/(2))\mathrm{Prob}_F(v_s \leqslant 1/2) \leqslant (v_b - 1/2)\mathrm{Prob}_F(v_s \leqslant 1/2)$. Thus $\hat{b}_1$ is optimal for $v_b \in [1/2,1]$. This proves optimality of $\hat{b}_1$. A similar argument holds for the seller's strategy, $\hat{s}_1$.

Stage $(n-1)$: Suppose that given any h^{n-1}, $(\hat{b}_{n-1}, \hat{s}_{n-1})$ constitutes an equilibrium.

Stage n: We show that $(\hat{b}_n, \hat{s}_n)$ is an equilibrium of the n^{th} stage game. Consider the buyer at stage n, with a true valuation $\hat{v}_b$, s.t., $\hat{v}_b \in [I/2^{n-1}, (I+1)/2^{n-1})$. We consider each of the possible cases next.

Suppose, to begin, that at stage n, $h^n = \{NT\}$, and so (F_n, G_n) are derived via Bayes' Rule, using the full history. From $(\hat{b}_{n-1}, \hat{s}_{n-1})$ buyer n knows that $\hat{b}_{n-1}(\hat{v}_b) = I/2^{n-1}$. Since $h_{n-1} = \{NT\}$, it follows that $\hat{s}_{n-1}(v_s) > I/2^{n-1}$, and so given $\hat{s}_{n-1}$, $l_f^n(\hat{v}_b) = I/2^{n-1}$ We show that $\hat{b}_n$ is optimal given the posterior beliefs, F_n, and the seller's strategy, $\hat{s}_n$. If $\hat{v}_b \in [2I/2^n), (2I+1)/2^n)$, then $\pi_n^b(b_n(v_b) \leqslant 2I/2^n,\ \hat{s}_n/F) = (v_b - 2I/2^n)\mathrm{Prob}_{F_n}(v_s \leqslant 2I/2^n) = 0$; $\pi_n^b((2I+1)/2^n > b_n(v_b) > 2I/2^n,\ \hat{s}_n/F) = (v_b - b)\mathrm{Prob}_{F_n}(v_s \leqslant 2I/2^n) = 0$; $\pi_n^b(b_n(v_b) > (2I+1)/2^n,\ \hat{s}_n/F_n) = (\hat{v}_b - b)\mathrm{Prob}_{F_n}(v_s \geqslant (21+1)/2^n) < 0$, since $\hat{v}_b < 2I + 1\}\{2^n$, and $\mathrm{Prob}_F(v_s \geqslant (21+1)/2^n) > 0$. Thus for $\hat{v}_b \in [2I/2^n, (2I+1)/2^n)$, $\hat{b}_n(v_b) = 2I/2^n$ is optimal. Similar arguments may be adduced to establish the optimality of $\hat{b}_n$ for other values of v_b.

Suppose next that $h^{n-1} = \{NT\}$, but $h_{n-1} = \{T\}$. If $\hat{v}_b \in [2I/2^n, (2I+1)/2^n)$, from $(\hat{b}_{n-1}, \hat{s}_{n-1})$ buyer n knows that $\hat{b}_{n-1}(\hat{v}_b) = I/2^{n-1}$. Since $h^{n-1} = \{NT\}$, and $h_{n-1} = \{T\}$, it follows that $\hat{s}_{n-1}(v_s) \leqslant I/2^{n-1}$, and given $\hat{s}_{n-1}$, $l_f^n(\hat{v}_b) = (I-1)/2^{n-1}$; and $u_f^n(\hat{v}_b) = I/2^{n-1}$. We wish to show that given this belief, $\hat{b}_n = \hat{b}_{n-1}$

$= 2I/2^n$ is optimal. $\pi_n^b(b_n(v_b) < 2I/2^n,\ \hat{s}_n/F) \leqslant v_b$. $\text{Prob}_{F_n}(v_s < I/2^{n-1}) = 0$, since $\text{Prob}_{F_n}(l_f^n(v_b)) = 0$; $\pi_n^b(b_n(v_b) = 2I/2^n,\ \hat{s}_n/F_n) = v_b - 2I/2^n)\ \text{Prob}_{F_n}(v_s \leqslant 2I/2^n) = (v_b - 2I/2^n)$; $\pi_n^b((2I+1)/2^n > b_n(v_b) > 2I/2^n,\ \hat{s}_n/F_n) = (v_b - (b + (2I/2^n))/(2))$. $\text{Prob}_{F_n}(v_s \leqslant 2I/2^n) < (v_b - 2I/2^n)$; $\pi_n^b(b > (2I+1)/2^n,\ \hat{s}_n/F_n) = (\hat{v}_b - (b + (2I+1)/2^n)/(2))$. $\text{Prob}_{F_n}(v_s \leqslant (2I+1)/2^n) < v_b - 2I/2^n$, since $\hat{b}_n(v_b) \geqslant (2I+1)/2^n$, and $\text{Prob}_{F_n}(v_s \leqslant (2I+1)/2^n) = 1$. Thus for $\hat{v}_b \in [2I/2^n,\ (2I+1)/2^n)$, $\hat{b}_n$ is optimal. Similar arguments may be adduced to establish the optimality of $\hat{b}_n$ for other v_b.

Consider off-equilibrium path histories next. Suppose h^n is commonly known to be inconsistent. Given that $G_n = G$ and $F_n = F$, $(\hat{b}_n,\ \hat{s}_n) = (\hat{b}_n,\ \hat{s}_n)$ is an equilibrium following arguments for stage 1. If h^n is some stage onward consistent, then arguments analogous to those above may be applied for $h^{n/k}$ consistent histories.

Suppose that h^n is s.t., $h^{n-2} = \{NT\}\ h_{n-1} = \{T\}$, but one and only one player, (say) the buyer, with valuation $\hat{v}_b$, knows that $\text{Prob}_{F_{n-1}}(v_s\colon \hat{b}_{n-1}(\hat{v}_b) \geqslant \hat{s}_{n-1}(v_s)) = 0$, i.e., h^n is inconsistent. Then by construction, $F_n = F_{n-1}$, and $\hat{b}_n(\hat{v}_b) = \hat{b}_{n-1}(\hat{v}_b)$. We demonstrate that $\hat{b}_{n-1}(\hat{v}_b)$ is optimal, given $F_g = F_{n-1}$, and $\hat{s}_n$. Suppose, for instance, that $\hat{v}_b \in [I/2^{n-1},\ I + 1/2^{n-1})$, and so $\hat{b}_{n-1}(\hat{v}_b) = I/2^{n-1}$. At stage n, by hypothesis this history is not commonly known to be inconsistent. In this case, for $\hat{v}_s \in [I/2_{n-1},\ I + 1/2_{n-1}$, history is consistent, and thus for such valuations, $\hat{s}_n(v_s) = \hat{s}_{n-1}$. Also note that for other seller valuations since history is inconsistent, by hypothesis, they also play, $s_n(v_s) = s_{n-1}(\hat{v}_s)$. Since $F_{n-1} = F_n$, the optimality of $\hat{b}_n$ then follows from the hypothesis that $\hat{b}_{n-1}$ is optimal. Analogous reasoning as in the case of consistent histories is sufficient to establish optimality of $\hat{s}_n$. An application of the principle of induction now completes the argument for existence. *Q.E.D.*

References

[1] P. Aghion, P. Bolton, C. Harris, B. Jullien, 1991, Optimal learning by experimentation, Review of Economic Studies 58, 4, 621–655.

[2] A. Ausubel and R. Deneckere, 1993, Efficient sequential bargaining, Review of Economic Studies 60, 3, 435–461.

[3] L. Blume and D. Easley, 1993, Rational learning and rational expectations, Mimeo, Cornell University.

[4] L. Blume and D. Easley, 1992, What has the rational learning literature taught us?, Mimeo, Cornell University.

[5] K. Chatterjee and W. Samuelson, 1983, Bargaining under asymmetric information, Operations Research 31, 835–51.

[6] C. d'Asperamont and P. Gerard-Varet, 1979, Incentives and incomplete information, J. of Pub. Econ 16, 25–45.

[7] S. Goyal, 1990, Essays on the economics of learning, Ph.D Thesis, Cornell University.

[8] S. Goyal, 1991, On the possibility of efficient bilateral trade, Mimeo, Econometric Institute, Erasmus Univerity, Rotterdam.

[9] S. Goyal, 1993, On the possibility of efficient bilateral trade, Mimeo, Econometric Institute, Erasmus University, Rotterdam.
[10] J. Harsanyi, 1967–68, Games with incomplete information played by Bayesian players, Management Science 14, 159–82, 320–334, 481–520.
[11] J. Jordan, 1991, Bayesian learning in normal form games, Games and Economic Behavior 3, 60–81.
[12] R. Myerson and M. Satterthwaite, 1983, Efficient mechanisms for bilateral trade, Journal of Economic Theory 29, 265–81.
[13] Symposium on non-cooperative bargaining, 1989, Journal of Economic Theory 48, 1.
[14] R. Wilson, 1987, Game theoretic analysis of trading processes, in: Advances in economic theory, ed., Truman F. Bewley, Cambridge University Press.
[15] R. Wilson, 1977, A bidding model of perfect competition, Review of Economic Studies 44, 511–518.
[16] A. Wolinsky, 1990, Information revelation in a market with pairwise meetings, Econometrica 58, 1–23.