

Commitment through Incomplete Information in a Simple Repeated Bargaining Game*

KLAUS M. SCHMIDT

*Department of Economics, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139*

Received July 11, 1989; revised March 10, 1992

This paper analyzes a finitely repeated bargaining game with asymmetric information. It gives a tight characterization of the equilibrium path and the equilibrium payoffs of all sequential equilibria satisfying a weak Markov property. The method used allows for arbitrarily many different types and provides an intuitive understanding of how “reputation” works. It is shown that the seller can use the incomplete information about his costs to credibly threaten never to accept a price lower than his highest possible costs before the very end of the game. The result holds for any discount factor of the seller greater than 0.5. *Journal of Economic Literature* Classification Numbers: C78, D83. © 1993 Academic Press, Inc.

1. INTRODUCTION

Bargaining Games with one-sided asymmetric information in which a buyer and a seller want to exchange one unit of a durable good are by now well understood.¹ Much less attention has been paid to the case of repeated bargaining where there is a perishable commodity or a service to be traded in every period. This situation differs substantially from the single trade bargaining case in two respects. On the one hand, the game does not end after a price offer has been accepted. Thus, if, e.g., the seller has private

* This paper is based on Chapter 2 of my Ph. D. thesis which was written within the European Doctoral Program at Bonn University and The London School of Economics. I thank in particular John Moore and Ariel Rubinstein for their advice and tireless encouragement. I also benefitted from helpful discussions with Helmut Bester, Georg Nöldeke, Monika Schnitzer, and Avner Shaked, and I am grateful to Drew Fudenberg, an associate editor, and two anonymous referees for their suggestions. Financial support by the Deutsche Forschungsgemeinschaft, SFB 303 at Bonn University, is gratefully acknowledged.

¹ See, e.g., Sobel and Takahashi [15], Fudenberg, Levine, and Tirole [3], or Gul, Sonnenschein, and Wilson [5].

information about his costs, he may be more reluctant to accept an offer if this reveals his costs to be low and gives him a smaller share of the surplus in future trades. This concern has been called "reputation effect" in other repeated games with asymmetric information. On the other hand, rejecting an offer is more costly than in the single trade bargaining case. The seller not only has to wait for another period, he also loses this period's trade irretrievably.

In this paper we show that the reputation effect strongly dominates the play. The seller will not accept any price offer lower than his highest possible costs before the very end of the game. Surprisingly, this result holds almost independently of the discount factors of the two players. Even if the buyer has a discount factor close to one and is much more patient than the seller (whose discount factor may be close to 0.5), she will not try to price discriminate in all but perhaps the last (bounded number of) periods. Thus, in the beginning of the game there is no transmission of information. This is in striking contrast to the equilibrium of the single trade bargaining game with repeated offers. There we have some price discrimination and transmission of information right from the beginning. Furthermore, according to the Coase conjecture the uninformed player (here the buyer) loses his bargaining power only if the discount factor of the seller is close to one, while discounting hardly matters in the repeated setting. Hence, the bargaining position of the uninformed party is even worse under repeated bargaining.

This paper is closely related to the literature on reputation effects started by Kreps and Wilson [8] and Milgrom and Roberts [12]. These authors have shown that a long-run player whose "type" (i.e., payoff function or strategy space) is private information can credibly threaten to behave as if he were another type, even if the prior probability of this type is very small. However, there is an important methodological difference. In their papers and in most of the subsequent literature a sequential equilibrium is calculated explicitly while the method used in our paper is to characterize equilibria. Without computing any equilibrium we can give a tight characterization of the equilibrium path and the equilibrium payoffs of all sequential equilibria satisfying a weak Markov property. The characterization approach has two main advantages. First, it allows for a more complex game structure with arbitrarily many different types and different discount factors for the seller and the buyer, while the calculation of an equilibrium is so complex that usually only two or three types can be considered. Second, our characterization follows a simple and intuitive line of reasoning which may provide a clearer understanding of the role of reputation.

The first who characterized rather than computed sequential equilibria were again Kreps, Milgrom, Roberts, and Wilson [7] in their work on the

finitely repeated prisoners' dilemma. Among the few other papers on equilibrium characterization in repeated games with incomplete information the one most closely related to ours is Fudenberg and Levine [2]. They allow for any perturbation of the information structure of the game which gives positive probability to a type who always plays a "commitment strategy," and they provide a lower bound for the equilibrium payoff of the informed player in any Nash equilibrium given that his discount factor is close to one. However, their result holds only if the uninformed player is completely myopic, or—equivalently—if the informed long-run player faces a sequence of uninformed short-run players, each of whom lives for one period only but knows the complete history of the game. We follow their approach in assuming that there exists a "commitment type" of the seller who accepts an offer if and only if it is greater or equal than the highest possible cost of the seller. All the other types of the seller differ only in their costs of production. Since this game has much more structure than the quite general class of games considered by Fudenberg and Levine [2] we can extend the analysis to the two long-run player case. In contrast to Fudenberg and Levine our lower bound for the equilibrium payoffs is tight for any discount factors bigger than 0.5 given that the horizon is long enough, and we can also characterize the equilibrium path.

A repeated bargaining game similar to the one discussed here has been analyzed by Hart and Tirole [6] who focus on the effects of long-term contracts and renegotiation. They consider the case of only two different types and common discount factors for both players. Both assumptions are crucial for their characterization argument to go through, and their proof follows a quite different line of reasoning.

The rest of the paper is organized as follows. Section 2 describes the repeated bargaining game and introduces the notation. In Sections 3 and 4 the main results are developed and discussed. Section 3 states some preliminary results which are used in Section 4 for the equilibrium characterization. The line of reasoning followed here gives a clear understanding of how "commitment through incomplete information" works. Section 5 deals with the existence problem. In Section 6 we conclude with some remarks on the transmission of information in equilibrium. It is shown that there is a close connection to the "ratchet effect" in a repeated principal-agent relationship. Then we analyze what is happening in the end of the game when the buyer starts to price discriminate. For the most simple case with only two different types of the seller we construct the (generically) unique equilibrium path. Finally, for the general case of arbitrarily many different types it is shown that information is not transmitted gradually (through successive skimming) as in a single trade bargaining game, but that equilibrium behavior is extremely complex with both players randomizing in the last periods.

2. THE BARGAINING GAME

A seller (S) can produce one unit of a good or of a service per time period with constant unit cost $c \leq 1$ which is his private information. The cost c may also be interpreted as his reservation value or his opportunity costs. In each period he wishes to sell the good to a buyer (B) with reservation price $b > 1$. Thus, trade is always efficient. The reservation price b is assumed to be common knowledge; e.g., the buyer has the outside option to buy the good elsewhere for the fixed price b .

Before the first stage game starts the "type" of the seller is drawn by nature out of the finite set $C = \{c^*, c^1, \dots, c^n\}$, with $1 = c^* > c^1 > \dots > c^n = 0$. It will be convenient to assume that all c^i are rational numbers. A cumulative distribution function $F(c)$ which is common knowledge represents the initial beliefs of the buyer about the type of the seller. $F(c)$ is the probability that his costs are less than or equal to c , and $\mu^i = F(c^i) - F(c^{i+1})$ denotes the prior probability that $c = c^i$, $i \in \{*, 1, \dots, n\}$. Everything else, including the complete history of the game, is common knowledge.

The following stage game G takes place in every period: The buyer makes a "take-it-or-leave-it" offer $p_t \in P$, where the price space $P \subset \mathbb{R}$ is specified below in this section. If the seller accepts p_t (A), he produces the good with cost c and trade takes place. If he rejects (R), there is no production, no trade, and no possibility of further bargaining in period t . Let G^T denote the stage game G repeated T times, $T < \infty$. Time is counted backwards ($t = T, T-1, \dots, 1$) for convenience.

The stage game payoffs are $p_t - c^i$ for the seller with cost c^i ($i = 1, \dots, n$) and $b - p_t$ for the buyer if the seller accepts p_t and 0 for both of them if he rejects. The seller with type c^* is called the "commitment type" who has a slightly different payoff function. Following Fudenberg and Levine [2] we assume that for him it is a dominant strategy in the repeated game to accept p_t if and only if $p_t \geq 1$. The idea is that there may be a small but positive probability that the seller is "committed" to accept price offers if and only if they are higher than a certain level which has been normalized to 1. We will show that the existence of such a type can be used by all the other types of the seller to credibly threaten to behave as if they were committed to this strategy as well. The assumption of a commitment type is a little bit stronger than necessary. In the following analysis we only need that there is a type of the seller who follows this strategy in every equilibrium. In Section 6 we show for the special case with only two different types of sellers with normal payoff functions that there is a generically unique sequential equilibrium in which the seller with cost $c = 1$ follows this strategy. However, if there are more than two types, we cannot rule out the possibility that a seller with cost $c = 1$ rejects $p_t = 1$.

This is why we assume that there exists a commitment type with positive probability.²

Let h_t be a specific history of the repeated game out of the set $H_t = (P \times \{A, R\})^{T-t}$ of all possible histories up to but not including period t . A (behavioral) strategy of the buyer is a sequence of functions, $\sigma_t^B: H_t \rightarrow \Sigma^B$, each of which maps the set of all possible histories H_t into the set Σ^B of all probability distributions over the price space P . The seller with cost c conditions his plan of action in period t not only on h_t , but also on the price offer p_t he got in this period. Hence, a (behavioral) strategy of the seller is a sequence of functions $\sigma_t^S(c): (H_t \times P) \rightarrow \Sigma^S$, where Σ^S is the set of all probability distributions over $\{A, R\}$. Finally, let $\mu_t^i(h_t)$, $i \in \{*, 1, \dots, n\}$, denote the updated probability the buyer assigns to the event that $c = c^i$ given that she observed the history h_t . The reference to history will be omitted if there is no ambiguity.

All players are risk neutral. The seller (buyer) discounts his (her) payoff in future periods with the discount factor $\delta < 1$ ($\beta < 1$). The expected payoffs of the buyer and of the seller with cost c^i ($i = 1, \dots, n$) from the beginning of period t for the rest of the game are given by

$$V_t^B = E_t \left\{ \sum_{j=0}^{t-1} \beta^j \cdot \pi_{t-j}(p_{t-j}) \cdot (b - p_{t-j}) \right\}, \quad (1)$$

$$V_t^S(c^i) = E_t \left\{ \sum_{j=0}^{t-1} \delta^j \cdot \pi_{t-j}(p_{t-j}) \cdot (p_{t-j} - c^i) \mid c^i \right\}, \quad (2)$$

where E_t is the expectation operator and $\pi_{t-j}(p_{t-j})$ denotes the probability that p_{t-j} is accepted in period $t-j$ given the history up to period $t-j$. To simplify notation the reference to history is omitted. The special case $\beta = 0$ is of particular interest. It is equivalent to the case of a sequence of

² An alternative specification is the following. Assume that the payoff function of type c^* in the stage game is given by

$$v_t(c^*, p_t) = \begin{cases} p_t - 1 - \theta \cdot \eta & \text{if } p_t \text{ is accepted} \\ (1 - \theta) \cdot \eta & \text{if } p_t \text{ is rejected,} \end{cases}$$

where $\eta > 0$ may be arbitrarily small and

$$\theta = \begin{cases} 0 & \text{if } p_t \geq 1 \\ 1 & \text{if } p_t < 1. \end{cases}$$

This payoff function is very close to the normal payoff function of a seller with cost of 1. It is easy to check that with this specification all our results go through.

buyers (B_t), each of whom knows the complete history of the game but lives for only one period and whose payoff function is

$$V_t^B = \pi_t(p_t) \cdot (b - p_t). \quad (3)$$

With arbitrarily many different types of the seller and an arbitrary distribution function $F(c)$ it is extremely complex to calculate any sequential equilibrium. We compute an equilibrium explicitly only for the very special case with two possible types of the seller in Section 6. For the general case, however, we characterize rather than calculate the set of equilibria. Here existence seems to be the main problem. In order to use a general existence theorem, the buyer's strategy space, i.e., the price space P , has to be finite.³ On the other hand, the price grid has to be sufficiently fine compared to the horizon (T) of the game. Since we are interested in the limit behavior of the set of equilibria if T goes to infinity, the parameters of the game have to be chosen in the following order. First, all parameters except T and P are fixed, then a large T is chosen, and then a sufficiently fine price grid P (as compared to T) is selected.

For any given T we construct the price grid as follows: Fix a real number $\eta > 0$ which may be arbitrarily small. Recall that we assumed the possible costs of the seller to be rational numbers. Thus, there exists an integer \bar{r} and integers k^1, \dots, k^n such that $c^i = k^i/\bar{r}$ and $1/\bar{r} < \eta$. Fix such an integer \bar{r} and define $r = \bar{r} \cdot 2^T$. Finally, fix an integer $l > r \cdot b \cdot T$. Then the price grid in the game G^T is chosen to be

$$P(\eta, T) = \{p^0, p^1, \dots, p^l\},$$

where $p^j = j/r$, $j = 0, 1, \dots, l$. Clearly, the buyer will never want to offer more than p^l , since $p^l > r \cdot b \cdot T/r = b \cdot T$, so the upper bound on prices is no restriction. If T increases the price grid becomes more and more close-meshed. Note that by construction the possible types of the seller are always on the grid.⁴

The appropriate solution concept for this game with incomplete information is the notion of sequential equilibrium (SE) as developed by Kreps and Wilson [9]. We provide a tight characterization of the equilibrium path and of the equilibrium payoffs of all sequential equilibria satisfying the following Weak Markov Property.

³ All our results (except existence) can be proved slightly more elegantly with a continuous price space, i.e., $P = \mathbb{R}$. See Schmidt [14].

⁴ This particular specification of the price grid will be convenient to work with in the proofs of the lemmata in the next section, but it is not necessary. We only have to require that the grid is fine enough for large T . For example, any grid $P' \supset P(\eta, T)$ would do the job as well. Nor is it necessary that the c^i are rational numbers lying on the grid, but it very much simplifies the exposition of the proofs.

DEFINITION. A sequential equilibrium of G^T satisfies the Weak Markov Property (WMP) if the seller's equilibrium strategy σ_t^S in period t is contingent only on his type, the current price offer p_t , and the public beliefs μ_t in period t .

How severe is the restriction to Markov Perfection? In Section 5 we show that a sequential equilibrium satisfying WMP always exists. Thus, if along the equilibrium path the equilibrium strategies are unique then the Weak Markov Property holds along the equilibrium path in all sequential equilibria. Generic uniqueness can be shown for the simple case with only two different types of sellers. For this case it is possible to construct the unique equilibrium path and to see that the requirement of the Weak Markov Property is no restriction (see Section 6.2). Furthermore, in the special case of a long-run seller facing a sequence of short-run buyers ($\beta = 0$) the restriction to equilibria satisfying the WMP is not necessary and all our results hold for all sequential equilibria. Thus the WMP has bite only if there are two long-run players with at least three different types of sellers. If there are multiple equilibrium outcomes Markov Perfection is a sensible refinement restricting the set of equilibria. It requires that the strategies of the players are as simple as is consistent with rationality. Furthermore it captures the idea that "bygones are bygones" more rigorously than does the notion of perfect equilibrium. That is, the outcome of a proper subform should not depend on the different paths by which this subform has been reached.⁵ Finally, in this specific game, the Weak Markov Property gives a tight characterization of the equilibrium path and the equilibrium payoffs, as will be shown in the next sections.

3. SOME PRELIMINARY RESULTS

The bargaining procedure is modelled in such a way that in each period all the bargaining power is on the side of the uninformed buyer. If she knew the cost of the seller she could extract all the surplus of trade by offering $p_t = c$ which would always be accepted in equilibrium. However, the seller's cost is private information and he may use this advantage to get some share of the surplus as an information rent. The following Lemmata 1 to 3 give a first characterization of the buyer's and the seller's behavior in equilibrium. Lemma 1 provides an upper bound for the prices offered in equilibrium and a lower bound for the equilibrium payoff of the buyer.

In Lemmata 1 and 2 it does not matter whether type c^* is a commitment type or whether he has a normal payoff function with cost $c = 1$. Recall

⁵ For a more detailed discussion of these points see Maskin and Tirole [11].

from the definition of the price grid $P(\eta, T)$ that for any arbitrarily small η the integer r has been chosen such that $2^T/r < \eta$.

LEMMA 1. *Consider any history up to but not including period t of G^T , and let c_t^m denote the maximum of all types to whom the buyer assigns positive probability after this history. If c^* has positive probability, then $c_t^m = 1$. On the equilibrium path of every SE satisfying the WMP all types $c \leq c_t^m$ will accept any price offer $p_t \geq c_t^m + 2^{t-1}/r$ and the buyer will offer $p_t \leq c_t^m + 2^{t-1}/r$. The expected equilibrium payoff of the buyer from period t onwards is bounded below by*

$$V_t^B(c_t^m) = \sum_{j=0}^{t-1} \beta^j \cdot \left(b - c_t^m - \frac{2^j}{r} \right) > \sum_{j=0}^{t-1} \beta^j \cdot (b - c_t^m) - \eta. \quad (4)$$

Proof. The proof is by induction. In the last period ($t = 1$) the buyer believes that $c \leq c_1^m$ with probability 1. If $p_1 > c_1^m$ is offered every $c \leq c_1^m$ strictly prefers to accept p_1 . Thus, $p_1 \leq c_1^m + 1/r$ will be offered in equilibrium. The buyer can guarantee herself a payoff of at least $V_1^B(c_1^m) = b - c_1^m - 1/r > b - c_1^m - \eta$ by just offering $p_1 = c_1^m + 1/r$.

Suppose the lemma is true for period t . In period $t+1$ the upper bound of the buyer's beliefs about c is c_{t+1}^m . Suppose $p_{t+1} \geq c_{t+1}^m + 2^t/r$ and there are some $c \leq c_{t+1}^m$ who reject p_{t+1} with positive probability. Take the maximum of these c and call it \bar{c} . If \bar{c} rejects he gets nothing in period $t+1$ and he has revealed that his costs are at most \bar{c} . Hence $c_{t-j}^m \leq \bar{c}$ ($j = 0, \dots, t-1$), because the updating of the beliefs has to be consistent in the equilibrium which is supposed to be played from this subform onwards. However, by the induction hypothesis the seller will not be offered more than $c_{t-j}^m + 2^{t-j-1}/r$ in the future. Thus, the overall payoff of type \bar{c} from rejecting p_{t+1} is at most

$$\delta \cdot \sum_{j=0}^{t-1} \delta^j \cdot \frac{2^{t-j-1}}{r} < \sum_{j=0}^{t-1} \frac{2^{t-j-1}}{r} = \frac{2^t - 1}{r}. \quad (5)$$

If he accepts p_{t+1} his payoff is at least $p_{t+1} - \bar{c} \geq 2^t/r$, so \bar{c} would have done better accepting. Therefore such a \bar{c} does not exist and every $c \leq c_{t+1}^m$ will accept $p_{t+1} \geq c_{t+1}^m + 2^t/r$.

Suppose the buyer offers $p_{t+1} > c_{t+1}^m + 2^t/r$ in equilibrium. Since any $p_{t+1} > c_{t+1}^m + 2^t/r$ will be accepted by all $c \leq c_{t+1}^m$, any such p_{t+1} will yield the same posterior beliefs about c in period t and thus, by the Weak Markov Property, the same equilibrium payoff for the buyer from period t onwards. Therefore, $p_{t+1} > c_{t+1}^m + 2^t/r$ cannot be an equilibrium price offer, and the buyer will offer $p_{t+1} \leq c_{t+1}^m + 2^t/r$. Suppose the buyer offers $p_j = c_{t+1}^m + 2^{j-1}/r$ in period j , $j = t+1, \dots, 1$. By the induction hypothesis all

these price offers will be accepted. Thus the buyer can guarantee herself at least

$$\begin{aligned}
 V_{t+1}^B(c_{t+1}^m) &= \sum_{j=0}^t \beta^j \cdot \left(b - c_{t+1}^m - \frac{2^j}{r} \right) \\
 &= \sum_{j=0}^t \beta^j \cdot (b - c_{t+1}^m) - \sum_{j=0}^t \beta^j \cdot \frac{2^j}{r} \\
 &> \sum_{j=0}^t \beta^j \cdot (b - c_{t+1}^m) - \frac{2^{t+1}}{r} \\
 &\geq \sum_{j=0}^t \beta^j \cdot (b - c_{t+1}^m) - \eta. \quad \text{Q.E.D.} \quad (6)
 \end{aligned}$$

Note that Lemma 1 rests on a backward induction argument which cannot be used if there is an infinite horizon. In fact, it is easy to construct an equilibrium of the infinitely repeated game where Lemma 1 does not hold. Note further that the Weak Markov Property has only been used in the induction step to show that the buyer will never offer $p_{t+1} > c_{t+1}^m + 2^t/r$. It is easy to see that in the case of a long-run seller facing a sequence of short-run buyers ($\beta = 0$) this follows trivially from the fact that the buyer in period $t+1$ is only interested in her payoff of the current period. Given that any $p_{t+1} \geq c_{t+1}^m + 2^t/r$ will be accepted with probability 1, a short-run buyer will never offer more than $c_{t+1}^m + 2^t/r$. Therefore, we do not need the WMP in this case. However, in the more general case of two long-run players ($\beta > 0$) the buyer might offer $p_{t+1} > c_{t+1}^m + 2^t/r$. Although this loses money in period $t+1$ he might be compensated in the future if this price offer leads to a continuation equilibrium which gives the buyer a higher payoff than he would have got otherwise.⁶ To deal with this problem the WMP is needed. All the following lemmata and propositions use Lemma 1. Thus, if not stated otherwise, from now on we mean by "equilibrium" a sequential equilibrium satisfying the WMP.

Lemma 2 says that it never pays for the seller to accept an offer significantly lower than his costs.

LEMMA 2. *If p_t is offered in period t , every seller with type $c \geq p_t + 2^{t-1}/r$ strictly prefers to reject p_t .*

Proof. Suppose there are some types $c \geq p_t + 2^{t-1}/r$ who accept p_t with positive probability. Take the maximum of these types and denote it by \bar{c} .

⁶ I am grateful to Drew Fudenberg and an anonymous referee who made me aware of this problem.

If the seller with cost \bar{c} accepts p_t , his payoff in period t is $p_t - \bar{c} \leq -2^{t-1}/r$ and he has revealed that his costs are $c \leq \bar{c}$. By Lemma 1 he will not be offered more than $p_j = \bar{c} + 2^{j-1}/r$ ($j = t-1, \dots, 1$) in the future. So the most he can gain in the future is

$$\sum_{j=1}^{t-1} \delta^j \cdot \frac{2^{j-1}}{r} < \sum_{j=1}^{t-1} \frac{2^{j-1}}{r} < \frac{2^{t-1}}{r}. \quad (7)$$

Therefore he cannot be compensated for his loss and strictly prefers to reject p_t . Q.E.D.

The special payoff function of the commitment type is exploited for the first time in the following Lemma 3, which states that $p_t \geq 1$ will always be accepted in equilibrium.

LEMMA 3. *If $p_t \geq 1$ is offered it will be accepted with probability 1 in any equilibrium. Furthermore, if $c_t^m = c^*$ the equilibrium payoff of the buyer in any subform starting in period t is bounded below by $\underline{V}^B = \sum_{j=0}^{t-1} \beta^j \cdot (b-1)$.*

Proof. The commitment type c^* accepts $p_t \geq 1$ by assumption. Suppose there is a c^i , $i \in \{1, \dots, n\}$, who rejects $p_t \geq 1$ in equilibrium with positive probability. Denote the maximal such type by \bar{c} . His payoff in period t is 0 and he has revealed that his costs are at most \bar{c} . From Lemma 1 we know that he will not be offered more than $\bar{c} + 2^{j-1}/r$ in periods $j = t-1, \dots, 1$. So his overall payoff is bounded above by $2^{t-1}/r < \eta$. However, by accepting he could have got $1 - \bar{c} \geq 1 - c^1 > \eta$, a contradiction. It follows immediately that the buyer can guarantee herself at least $\underline{V}^B = \sum_{j=0}^{t-1} \beta^j \cdot (b-1)$ by just offering $p_t = 1$ in every period. Q.E.D.

Lemma 3 implies that if the commitment type of the seller has positive probability, then in any equilibrium there exists with positive probability a history in which p_t is accepted if and only if $p_t \geq 1$. This is the crucial ingredient which is necessary to establish our main results. Note that by Lemmata 1 and 2 a seller with cost $c = 1$ accepts any $p_t \geq 1 + \eta$ and rejects any $p_t \leq 1 - \eta$, where $\eta > 0$ may be arbitrarily small. Furthermore, for a continuous price grid it can be shown that a seller with $c = 1$ accepts p_t if $p_t > 1$ and rejects if $p_t < 1$. However, if $p_t = 1$ is offered he is indifferent and we cannot show that he is going to accept p_t in equilibrium. This is why the commitment type has to be assumed.

Lemma 4 focuses on such a history, where p_t has been accepted if and only if $p_t \geq 1$, and deals with the following question: Suppose the buyer offers $p_t < 1$ in equilibrium. Is there a lower bound $\varepsilon > 0$ independent of t , such that the probability of acceptance of p_t , denoted by $\pi_t(p_t)$, must be at least ε ? Such a lower bound clearly exists in the special case of a

sequence of short-run buyers, each of whom lives for one period only ($\beta = 0$). By Lemma 3 a lower bound for the equilibrium payoff of a short-run buyer is given by $\underline{V}^B = b - 1$. Thus, if she offers $p_t < 1$ in equilibrium it must be true that $\pi_t(p_t) \cdot (b - p_t) \geq b - 1$, or

$$\pi_t(p_t) \geq \frac{b-1}{b-p_t} \geq \frac{b-1}{b} > 0. \quad (8)$$

However, in the general case of a long-run buyer ($\beta > 0$) this need no longer be true. A long-run buyer is concerned about both her present and her future payoff, so she may trade off losses today against gains tomorrow. In particular, it might pay for her to invest in gathering information in the beginning of the game, i.e., to "test" the type of the seller with price offers lower than 1 even if the probability of acceptance is very small. Lemma 4 says that although for any $p_t < 1$ the probability of acceptance may be arbitrarily small it cannot be arbitrarily small for arbitrarily many price offers.

LEMMA 4. Assume $\mu^* > 0$. Fix an integer M ,

$$M \geq N = \frac{\ln(1-\beta) + \ln(b-1) - \ln b}{\ln \beta} > 0, \quad (9)$$

and define a positive number ε ,

$$\varepsilon = \frac{(1-\beta) \cdot (b-1)}{b} - \beta^M > 0. \quad (10)$$

Consider a history h_t of any equilibrium in which all price offers p_j ($j = T, \dots, t$) have been accepted if and only if $p_j \geq 1$. If there have been M price offers $p_j < 1$ along such a history, at least one of them must have had a probability of acceptance of at least ε .

The proof is lengthy and relegated to the Appendix but the underlying intuition can be outlined briefly. By Lemma 3 we know that $\underline{V}_t^B = \sum_{j=0}^{t-1} \beta^j \cdot (b-1)$ is a lower bound for the equilibrium payoff of the buyer. Thus, offering $p_t < 1$ must yield an expected payoff of at least \underline{V}_t^B . On the other hand $\bar{V}_t^B = \sum_{j=0}^{t-1} \beta^j \cdot b$ clearly is an upper bound for her equilibrium payoff after any history h_t because the seller will guarantee himself at least 0. Now suppose the buyer "invests" in screening the different types of sellers by making M price offers smaller than 1 although she believes that each of them will be rejected with a probability higher than $1 - \varepsilon$. The possible return of such an investment is bounded above by \bar{V}_t^B . However, if the buyer discounts her future payoffs, the return from an investment

may not be delayed too far to the future. We have constructed ε and M such that if the buyer makes M price offers smaller than 1, each of them having a probability of acceptance of less than ε , then her highest possible expected payoff is smaller than V_i^B , a contradiction.

4. EQUILIBRIUM CHARACTERIZATION

The following proposition states that if all offers $p_i < 1$ have been rejected then the buyer will offer prices lower than 1 at most K times in equilibrium, where K is a finite number depending on β , b , and μ^* but independent of T . Thereafter she will offer $p_i = 1$ until the game ends. The intuition behind this result is simple and gives a clear understanding of how the building of reputation works: If the buyer repeatedly offers prices smaller than 1 she ultimately expects that some of these prices will be accepted with positive probability (this is the role of Lemma 4). However, only a seller with $c < 1$ may accept $p_i < 1$. Hence, if $p_i < 1$ is rejected, the updated probability that $c = c^*$ increases. If the buyer has expected p_i to be accepted with a probability higher than $\varepsilon > 0$, then the updated probability that $c = c^*$ after p_i has been rejected has to increase by an amount bounded away from 0. But the updated probability that $c = c^*$ is bounded above by 1. Therefore this can happen only in a bounded number of periods. In other words, if the seller has rejected price offers lower than 1 often enough, then the updated probability that he is the commitment type becomes so big that the buyer will no longer try to price discriminate. This argument is made precise in the proof of the following proposition.⁷

PROPOSITION 1. *Let $0 \leq \beta < 1$, $0 \leq \delta \leq 1$, $b > 1$, $\mu^* > 0$, and let $[x]$ denote the integer part of x . Consider any equilibrium of G^T and any history consistent with this equilibrium along which the seller has accepted p if and only if $p \geq 1$. Along such a history there cannot be more than K price offers $p < 1$, where*

$$K = K(\beta, b, \mu^*) = M \cdot \left(\left\lceil \frac{\ln \mu^*}{\ln((\beta b - \beta + 1)/b + \beta^M)} \right\rceil + 1 \right), \quad (11)$$

with

$$M = [N] + 1 = \left\lceil \frac{\ln(1 - \beta) + \ln(b - 1) - \ln b}{\ln \beta} \right\rceil + 1 \geq 1. \quad (12)$$

⁷ The part on Bayesian updating in the proof of Proposition 1 is similar to the proof of Lemma 1 in Fudenberg and Levine [2]. However, in the setting of our model the argument is simpler and has been developed independently.

Proof. Note first that if $\mu^* > 0$ then with positive probability there exists a history h^* , along which the seller accepts p iff $p \geq 1$, in every equilibrium of G^T . Assume there is an equilibrium for which the property does not hold, i.e., there is a history h^* in which the seller has accepted p_t iff $p_t \geq 1$ ($t = T, \dots, 1$) and in which the buyer has offered $p_t < 1$ more often than K times. We will show that this leads to a contradiction. Recall that $\mu_t^* = \mu_t^*(h_t)$ denotes the updated probability that $c = c^*$ given history h_t . By Lemma 3 we know that prices bigger or equal than 1 will be accepted by all types, so the probability assessment remains unchanged if $p_t \geq 1$ has been accepted. Prices lower than 1 will be rejected for sure by the seller with type c^* . Therefore, along the history h_t^* the updated probability $\mu_t^*(h_t^*)$ cannot decrease, where h_t^* is the history h^* truncated at period t .

Divide the K (or more) price offers lower than 1 to successive groups of M , $M = [N] + 1 > N = (\ln(1 - \beta) + \ln(b - 1) - \ln b) / \ln \beta \geq 0$ and consider the first group of M of them. By Lemma 4 we know that at least one of them (call it p_{τ_1}) must have a probability of acceptance of at least $\varepsilon = (1 - \beta)(b - 1) / (b - \beta^M) > 0$. Therefore the updated probability that $c = c^*$ after p_{τ_1} has been rejected is

$$\begin{aligned} \mu_{\tau_1-1}^* &= \text{Prob}(c = c^* \mid p_{\tau_1} \text{ was rejected in } \tau_1) \\ &= \frac{\text{Prob}(c = c^* \text{ and } p_{\tau_1} \text{ was rejected in } \tau_1)}{\text{Prob}(p_{\tau_1} \text{ was rejected in } \tau_1)} \\ &= \frac{\mu_{\tau_1}^*}{1 - \pi_{\tau_1}(p_{\tau_1})} \geq \frac{\mu_{\tau_1}^*}{1 - \varepsilon} \geq \frac{\mu^*}{1 - \varepsilon}. \end{aligned} \quad (13)$$

Now take the second group of M price offers lower than 1. Again at least one of them (p_{τ_2}) has a probability of acceptance higher than ε . This gives

$$\mu_{\tau_2-1}^* \geq \frac{\mu_{\tau_2}^*}{1 - \varepsilon} \geq \frac{\mu_{\tau_1-1}^*}{1 - \varepsilon} \geq \frac{\mu^*}{(1 - \varepsilon)^2}. \quad (14)$$

Consider the n th group of M price offers lower than 1 and denote by p_{τ_n} one of the offers with a probability of acceptance higher than ε . By induction the updated probability ($\mu_{\tau_n-1}^*$) that $c = c^*$ after p_{τ_n} has been rejected is bounded below:

$$\mu_{\tau_n-1}^* \geq \frac{\mu^*}{(1 - \varepsilon)^n}. \quad (15)$$

However, $\mu_{\tau_n-1}^* \leq 1$. Therefore there is an upper bound for n , namely

$$n \leq \frac{\ln \mu^*}{\ln(1 - \varepsilon)}. \quad (16)$$

Substituting for ε gives

$$n \leq \frac{\ln \mu^*}{\ln((\beta b - \beta + 1)/b + \beta^M)}. \quad (17)$$

Thus, an upper bound for the number of periods in which price offers smaller than 1 are offered and rejected on the equilibrium path is given by

$$K(\beta, b, \mu^*) = M \cdot \left[\frac{\ln \mu^*}{\ln((\beta b - \beta + 1)/b + \beta^M)} + 1 \right]. \quad (18)$$

Note that K is increasing in M , so choosing $M = [N] + 1$ minimizes K . Thus we have a contradiction to the assumption that there are more than K such price offers. Q.E.D.

It is evident from the proof how K depends on μ^* , b , and β . The bigger μ^* the less updating is needed to hit the point where $\mu_{t_n-1}^*$ would become higher than 1, and the smaller is K . Furthermore K is a decreasing function of b (note that the numerator and the denominator are both negative). The reason is that the lower bound for the equilibrium payoff of the buyer is increasing with b . Thus, if she risks price offers smaller than 1, the possible loss is bigger and therefore the probability of acceptance must be higher. This gives a bigger ε and a smaller K . Finally, K is an increasing function of β . If the buyer is more patient, then future gains become more important and she may offer prices lower than 1 with a smaller probability of acceptance in order to invest in gathering information. This effect increases the number of periods in which prices lower than 1 can be offered and rejected. On the other hand, if $\beta \rightarrow 0$ it is easy to see that $N \rightarrow 0$ and

$$K(0, b, \mu^*) = \left[-\frac{\ln \mu^*}{\ln b} \right]. \quad (19)$$

This is the same upper bound as the one derived by Fudenberg and Levine [2] for the special case of an informed long-run player facing a sequence of uninformed short-run players.⁸

⁸ Fudenberg and Levine [2] argued that their results do not apply to the two long-run player case because an uninformed long-run player may take an action in equilibrium which is not a short-run best response to the "commitment strategy" of the informed player, even if the probability that the "commitment strategy" will be played is arbitrarily close to 1. This is certainly true. However, our argument here is that this may not happen arbitrarily often if the uninformed player discounts her future payoffs and if her future payoffs are bounded above (see Lemma 4), which gives the positive lower bound for the probability of acceptance in at least some periods. This argument can be established more generally for the class of all two long-run player games with "conflicting interests" which is discussed in Schmidt [13]. Note, however, that the bargaining game considered here is not a game of "conflicting interests" as defined there.

Proposition 1 has an immediate implication for the equilibrium payoff of the seller. Suppose a seller with type $c < 1$ decides to build up a "highest cost" reputation and mimics the commitment type. The worst that can happen to him is that K price offers lower than 1 are offered and have to be rejected and that this happens in the first K periods of the game. In this case he would get 0 in the first K periods and $1 - c$ thereafter.

COROLLARY 1. *In the game G^T a lower bound for the equilibrium payoff of the seller with type c is given by*

$$\begin{aligned} V^S(c) &= \sum_{j=0}^{K-1} \delta^j \cdot 0 + \sum_{j=K}^{T-1} \delta^j \cdot (1-c) \\ &= (1-c) \cdot \delta^K \cdot \frac{1 - \delta^{T-K+1}}{1 - \delta}. \end{aligned} \quad (20)$$

However, this lower bound need not be very informative. Depending on β , b , and μ^* it may take quite a long time until the buyer starts to offer $p = 1$. Therefore, this lower bound gives a tight characterization of equilibrium payoffs only if the seller is very patient (δ is close to 1) and if the time horizon is sufficiently long.

Consider again the interpretation of Proposition 1. Suppose a seller with $c < 1$ mimics the commitment type. It follows from Proposition 1 that there will be at most $K(\beta, b, \mu^*)$ price offers lower than 1. If the seller rejects $p_t < 1$ often enough the probability that he is of type c^* increases until it becomes so big that the buyer will no longer try to price discriminate. In this sense the seller may invest repeatedly in the reputation of being the commitment type (i.e., in the updated probability that he is of type c^*). The costs of this investment are the missed opportunities of trade. The returns are the price offers of 1, which he will get for sure after K such rejections. Thus, Proposition 1 tells us that in principle every type of seller can build up this reputation. The question arises whether there are equilibria of the game with some types of the seller investing into this reputation in the beginning of the game in order to get better price offers thereafter. The answer is no, for the paradoxical reason that the incentives to do so are too big. Every type of seller would want to build up this reputation as long as the end of the game is not too close. Hence, every type would reject price offers lower than 1, but anticipating this the buyer will offer $p_t = 1$ in all periods before the very end of the game. Therefore the type of the seller will not be tested in the beginning of the game, the probability that he is of type c^* will remain at μ^* , and there is no reputation building in all but perhaps some of the last finite number of periods.⁹

⁹ Hart and Tirole [6, Proposition 3] provide a similar but more restricted result.

PROPOSITION 2. *Let β, δ, b , and μ^* satisfy $0 \leq \beta < 1$, $0.5 < \delta \leq 1$, $b > 1$, and $\mu^* > 0$. Define $L = L(\delta)$ as the smallest integer such that $\sum_{j=1}^L \delta^j > 1$ and choose η such that*

$$\sum_{j=1}^L \delta^j \geq 1 + \frac{\eta}{1-c^1}. \quad (21)$$

If $T > K(\beta, b, \mu^) \cdot L(\delta)$, then along the equilibrium path of every equilibrium of G^T a price of 1 is offered in at least the $T - K \cdot L$ first periods.*

Proof. Take any equilibrium and fix a history h^* in which the seller has accepted p iff $p \geq 1$. From Proposition 1 it is known that after at most K rejections of price offers lower than 1 the buyer will offer a price of 1 from then on. Suppose there have been $K-1$ such rejections up to period τ_1 and that there are more than L periods left ($\tau_1 > L$). If $p_{\tau_1} < 1$ is offered and rejected the payoff of the seller from period τ_1 onwards is

$$V_{\tau_1}(c \mid p_{\tau_1}, R) = 0 + \sum_{j=1}^{\tau_1-1} \delta^j \cdot (1-c). \quad (22)$$

Using Lemma 1 the highest payoff he can expect if he accepts p_{τ_1} is given by

$$\bar{V}_{\tau_1}(c \mid p_{\tau_1}, A) = p_{\tau_1} - c + \sum_{j=1}^{\tau_1-1} \delta^j \cdot \left(\bar{c} + \frac{2^{\tau_1-j-1}}{r} - c \right), \quad (23)$$

where \bar{c} is the maximum of all types who would accept p_{τ_1} with positive probability. For type \bar{c} it must be true that

$$p_{\tau_1} - \bar{c} + \sum_{j=1}^{\tau_1-1} \delta^j \cdot \frac{2^{\tau_1-j-1}}{r} \geq (1-\bar{c}) \cdot \sum_{j=1}^{\tau_1-1} \delta^j. \quad (24)$$

But this is a contradiction because

$$\begin{aligned} p_{\tau_1} - \bar{c} + \sum_{j=1}^{\tau_1-1} \delta^j \cdot \frac{2^{\tau_1-j-1}}{r} &< 1 - \bar{c} + \frac{2^{\tau_1-1}}{r} < 1 - \bar{c} + \eta \\ &= (1-\bar{c}) \cdot \left(1 + \frac{\eta}{1-\bar{c}} \right) \leq (1-\bar{c}) \cdot \left(1 + \frac{\eta}{1-c^1} \right) \leq (1-\bar{c}) \cdot \sum_{j=1}^{\tau_1-1} \delta^j. \end{aligned} \quad (25)$$

Therefore there is no such \bar{c} and all types will reject any $p_{\tau_1} < 1$.

Suppose the buyer would offer $p_{\tau_1} < 1$ in equilibrium nevertheless. Then his payoff from τ_1 onwards is given by

$$V_{\tau_1}^B(p_{\tau_1}) = 0 + \beta \cdot V_{\tau_1-1}^B(p_{\tau_1}, R). \quad (26)$$

By Lemma 3, a price offer \tilde{p}_{τ_1} with $1 \leq \tilde{p}_{\tau_1} < b$ would have been accepted with probability 1 yielding a payoff

$$V_{\tau_1}^B(\tilde{p}_{\tau_1}) = b - \tilde{p}_{\tau_1} + \beta \cdot V_{\tau_1-1}^B(\tilde{p}_{\tau_1}, A). \quad (27)$$

However, nothing is learned about the type of the seller, neither after p_{τ_1} has been rejected nor after \tilde{p}_{τ_1} has been accepted, and the probability assessment of the buyer is the same after both price offers. Thus, by the WMP, her equilibrium payoff in both subforms has to be the same and $V_{\tau_1}^B(\tilde{p}_{\tau_1}) > V_{\tau_1}^B(p_{\tau_1})$, a contradiction. Therefore, the buyer will not offer $p_{\tau_1} < 1$ in equilibrium.

We can go on by backward induction: Suppose there have been $K-2$ rejections of price offers lower than 1 up to a period τ_2 . If $\tau_2 > 2 \cdot L$, then every seller will get price offers of 1 at least in the next L periods if he rejects, and because of the above argument every type will do so. But then the buyer will offer a price of 1 in this period already. The proof goes on by induction until there have been $K-K=0$ rejections of price offers lower than 1. Q.E.D.

Note that the WMP is necessary (as in Lemma 1) to deal with the problem of multiple equilibria giving different continuation payoffs to the buyer in identical subforms. However, if the buyer is only interested in her payoff of the current period ($\beta=0$) this restriction is not necessary and Proposition 2 holds for all sequential equilibria.

Proposition 2 is rather insensitive to variations of δ . If $\delta \in [0.62, 1)$ then $L(\delta) = 2$, i.e., independent of the patience of the seller a price of 1 will be offered in the first $T - 2 \cdot K$ periods of the game. For large T this gives a tight lower bound for the equilibrium payoff of the seller independent of δ . This is quite surprising given the earlier lower bound $\underline{V}^S(c)$ of Corollary 1 which is tight only if δ is very close to 1.

Suppose that before the beginning of period T every type of the seller could publicly commit never to accept an offer smaller than 1. In this commitment version of G^T the unique equilibrium, which will be called the "commitment equilibrium," prescribes that the buyer always offers $p_t = 1$ which the seller will always accept. Given this definition we can state our main result, which follows directly from Propositions 1 and 2:

THEOREM 1. *Let β, δ, b , and μ^* satisfy $0 \leq \beta < 1$, $0.5 \leq \delta \leq 1$, $b > 1$, and $\mu^* > 0$. Consider a sequence of games G^T . If T goes to infinity the equilibrium paths of all equilibria of G^T converge to the commitment equilibrium path. That is, there are finite numbers $K = K(\beta, b, \mu^*)$ and $L = L(\delta)$ independent of T such that for all T and $t \geq T - K \cdot L$ the buyer offers $p_t = 1$. If*

T goes to infinity the equilibrium payoffs of the buyer and the seller converge to

$$V_{\infty}^B = \frac{b-1}{1-\beta} \quad \text{and} \quad V_{\infty}^S = \frac{1-c}{1-\delta}. \quad (28)$$

Theorem 1 gives a tight characterization of the equilibrium path and the equilibrium payoffs for large T . Perhaps surprisingly, this result holds even if the buyer is much more patient than the seller, as long as $\beta < 1$ and $\delta > 0.5$. Intuitively one might have expected that a patient uninformed buyer who deals with a relatively impatient informed seller will try to screen the different types of sellers with price offers lower than 1 in early periods in order to get more of the surplus of trade thereafter. Theorem 1 states that this is not the case. The reason is that as long as $\delta > 0.5$ every type of seller has a strong incentive to build up the reputation of being the commitment type. Therefore everybody will reject $p_t < 1$ in all but the last finite number of periods. But then there is no way to screen the different types and the buyer does better by offering 1, even if she is very patient. Note, however, that there is no actual “reputation building” on the equilibrium path, i.e., the updated probability that $c = c^*$ remains unchanged until the very end of the game. It is the mere possibility to build up a reputation which prevents the buyer from offering prices lower than 1. Thus it is perhaps more accurate not to talk about “reputation building” but to say that the seller can use the incomplete information about his type to credibly threaten to behave as if he were the commitment type. This is why we call this effect “commitment through incomplete information.”

5. EXISTENCE

Up to now all sequential equilibria satisfying the WMP have been characterized without showing that there exists at least one such equilibrium. For very simple cases (e.g., if there are only two different types of the seller and if there is no discounting, see Subsection 6.2) it is possible to show existence (and generic uniqueness) by constructing the equilibrium path. But for arbitrarily many different types and different discount factors this task is hopelessly complicated and a general existence theorem has to be used. By Proposition 1 of Kreps and Wilson [9] there exists a sequential equilibrium in G^T . However, if $\beta > 0$ we have to show that there is a sequential equilibrium of G^T satisfying the Weak Markov Property. Recall that the WMP restricts only the equilibrium strategy of the informed party

(the seller) but not of the uninformed player. The WMP is clearly satisfied in a “strong Markov equilibrium,” i.e., in a sequential equilibrium in which the (behavioral) strategies of both players depend only on the current beliefs and on the current price offer p_t , but it is well known that strong Markov equilibria do not in general exist.¹⁰ However, in a recent paper Maskin and Tirole [11] have shown existence for a slightly weaker equilibrium concept which they call “Markov Perfect Bayesian Equilibrium” (MPBE). We can state the following corollary to their Proposition 9.

COROLLARY 2. *There exists a Markov Perfect Bayesian Equilibrium in G^T . Every MPBE is a sequential equilibrium.*

Although the definition of a MPBE is quite complex and not provided here it is easy to show that in G^T every MPBE satisfies the Weak Markov Property. Thus there always exists at least one equilibrium which can be characterized as in Section 4.

6. TRANSMISSION OF INFORMATION IN EQUILIBRIUM

The equilibrium characterization of Section 4 has shown that the seller cannot be induced to transmit any information about his costs to the buyer in the beginning of the repeated relationship. This confirms the so-called “ratchet effect.” The literature on the ratchet effect offers a slightly different interpretation of our results which focuses on the information rents of the informed party and which may be useful to interpret our model from a different perspective. Thereafter we show how information is transmitted in the end of the game when the buyer starts to price discriminate. We conclude with some remarks on a possible extension of the model.

6.1. *The Ratchet Effect Revisited*

Laffont and Tirole [10] show that an agent with (privately known) high productivity has a strong incentive to pool with the low productivity type in order to hide his ability in the first period. The reason is that the principal cannot commit in the beginning of period 1 to an incentive scheme for period 2. If she learns the agent’s productivity in the first period she will offer an incentive scheme in the second period which does not give any

¹⁰ See, e.g., Kreps and Wilson [8] or Fudenberg and Tirole [4]. In their non-existence examples it is always the uninformed party who has to condition her strategy on past actions in equilibrium.

information rent to the agent. The agent could only be induced to reveal his productivity if the principal would offer the second period's information rent already in the first period. But then the low productivity agent could mimic the good type, take the money, and refuse to work (taking his outside option) in the second period. Thus, if the discount factor of the agent is not too small, there will be "substantial pooling" in all equilibria of this game.¹¹

Our model extends this idea to a repeated bargaining model with an arbitrarily long horizon. The intuition is very much the same. A low cost seller has a strong incentive not to reveal his costs too early because he will not be offered more than his costs thereafter. To compensate him for the foregone information rent the buyer would have to offer a very high price in the beginning ($p > 1$). However, this offer would also be accepted by the high cost types who can just refuse to sell for the low price offers thereafter. The crucial point is that once a seller has revealed his type he will not get any information rent thereafter. Together with the monotonicity of the seller's preferences in both his costs and the price offers this implies that there has to be complete pooling in all but the last periods.

6.2. Information Transmission with Two Types

In the end of the game, however, there is at least some transmission of information. It is possible to show what happens in the end of the game for the most simple case with only two different types of the seller and no discounting. Since we show existence in this case by constructing the equilibrium we can drop the assumption of a finite price grid and set $P = \mathbb{R}$. Furthermore, in this case we can show that a seller with cost of 1 will accept p_t if and only if $p_t \geq 1$ in equilibrium. So the behavior of the commitment type is determined endogenously. The (generically) unique equilibrium path of this case is given in Proposition 3.

PROPOSITION 3. *Let $c \in \{0, 1\}$, $\delta = \beta = 1$, $b > 1$, and $\mu^* = \mu > 0$. Consider the generic case where there is no integer m , such that $\mu = 1/b^m$. Then there exists a sequential equilibrium and every SE has the following "on the equilibrium path" strategies and value functions:*

Strategy of S with $c = 1$. Accept p_t iff $p_t \geq 1$.

Strategy of S with $c = 0$. If $p_t < 1$, reject p_t with probability $x_t = \min\{(\mu_t/(1 - \mu_t)) \cdot (b^{t-1} - 1), 1\}$, otherwise accept.

¹¹ In the language of principle agent theory the difference to a one period setting is that incentive constraints are binding in both directions. Even in the two period case the structure of the equilibria of this game with a continuum of types is very complicated. See Laffont and Tirole [10] for a characterization.

Strategy of B. Offer $p_t = 0$ if $\mu_t \in [0, 1/b')$ and $p_t = 1$ if $\mu_t \in (1/b', 1]$. If $\mu_t = 1/b'$, offer $p_t = 1$ with probability p_{t+1} and $p_t = 0$ with probability $(1 - p_{t+1})$.

$$V_t^S(c = 1) = 0$$

$$V_t^S(c = 0) = \begin{cases} t - n & \text{if } \exists n, \text{ such that } \frac{1}{b^n} \geq \mu_t > \frac{1}{b^{n+1}} \text{ and } n < t \\ p_{t+1} & \text{if } \exists n, \text{ such that } \frac{1}{b^n} = \mu_t \text{ and } n = t \\ 0 & \text{if } \exists n, \text{ such that } \frac{1}{b^n} > \mu_t \geq \frac{1}{b^{n+1}} \text{ and } n \geq t \end{cases}$$

$$V_t^B = \begin{cases} t \cdot b - (t - n) - \mu_t \cdot \sum_{j=1}^t b^j & \text{if } \exists n, \text{ such that } \frac{1}{b^n} \geq \mu_t > \frac{1}{b^{n+1}} \text{ and } n < t \\ t \cdot b - \mu_t \cdot \sum_{j=1}^t b^j & \text{if } \exists n, \text{ such that } \frac{1}{b^n} \geq \mu_t > \frac{1}{b^{n+1}} \text{ and } n \geq t. \end{cases}$$

The proof of the proposition is by induction and is left to the reader. Note that the equilibrium strategies satisfy the WMP. However, the equilibrium is not strong Markov because the buyer's action depends on his last period's price offer if $\mu_t = 1/b'$.

The pattern of the equilibrium path is familiar from the equilibrium path of the chain store game analyzed by Kreps and Wilson [8]. It is graphically illustrated in Fig. 1. In the beginning of the game, i.e., as long as $1/b' < \mu$, the buyer always offers 1, which is accepted by both types of sellers, and her probability assessment remains unchanged. If $t = n$, where n is the biggest integer such that $\mu < 1/b^n$, the buyer starts offering 0, which is rejected for sure by $c = 1$. The seller with type $c = 0$ randomizes. If he accepts $p_n = 0$, he has revealed his type and will accept $p_{t-j} = 0$ from then on with probability 1. If he rejects p_n , the buyer updates her probability assessment that $c = 1$ such that she is just indifferent between offering $p_{n-1} = 0$ and $p_{n-1} = 1$. However, in equilibrium she will offer $p_{n-1} = 0$ with probability 1. Otherwise the seller would have strictly preferred to reject $p_n = 0$. In period $n - 1$ the seller with cost of 0 will again randomize, but he now accepts with a higher probability, and so on. Only in the last period the seller with $c = 0$ will accept $p_1 = 0$ with probability 1 given that he has rejected all other price offers smaller than 1 before.




FIG. 1. Evolution of beliefs along the equilibrium path.

Note that the expected payoff of the low cost seller after period n is 0. Nevertheless he has to randomize between accepting and rejecting in order to prevent the buyer from offering a price of 0 already before period n . Reputation building, in the sense that μ_t increases, might only occur after period n when the expected payoff of the seller is 0 anyway. Thus, the seller does not actually gain from building up a reputation. It is also misleading to say that he will cash in his reputation when he finally accepts $p_{n+1} = 0$ because this just gives a payoff of 0. The important point is that if the seller were offered prices smaller than 1 before period n then it would pay for him to build up a reputation by rejecting them. Thus it is only the possibility to build up a reputation which prevents the buyer from testing his type.

6.3. No Successive Skimming

If there are more than two different types of the seller the analysis of the end of the game is much more complex. However, we can get some insights by comparing the repeated bargaining game with a single trade bargaining game with repeated offers as analyzed by Fudenberg, Levine, and Tirole [3] or Gul, Sonnenschein, and Wilson [5] and the literature on the Coase conjecture. In these games an uninformed seller makes a sequence of take-it-or-leave-it price offers to an informed buyer whose reservation value v is private information. Once an offer has been accepted trade takes place and the game ends, while in our repeated bargaining game there may be trade in every period. A common feature of all single trade bargaining games is the "successive skimming property" which says that after any history h_t and for any price offer p_t there is a cut-off type $\tilde{c}(p_t, h_t)$, such that p_t is

accepted if and only if $c \leq \tilde{c}(p_t, h_t)$.¹² Successive skimming implies that every type of the seller plays according to a pure strategy in equilibrium. Furthermore, information is transmitted gradually in the sense that the higher the cost of the seller the later he will accept a price offer smaller than 1. Thus, the successive skimming property very much facilitates the updating of the beliefs (a tail of the density function is cut off after every period) and the construction of an equilibrium.

In a repeated bargaining game, however, this property does not hold.¹³ This is the main analytical difference to single trade bargaining games with repeated offers and the main reason why we cannot say more about information transmission in the end of G^T . Of course, the equilibrium outcomes are different as well. In the single trade bargaining game there is some price discrimination in the beginning of the game, the extent of which depends on the discount factors. In the repeated bargaining game, however, price discrimination may only occur in the end of the game. This is why discounting plays a much weaker role here.

6.4. Two-Sided Uncertainty

An interesting extension of the model would be to allow for two-sided uncertainty, e.g., that the valuation of the buyer is private information as well. As long as the lowest possible reservation value of the buyer strictly exceeds the highest possible cost of the seller the results of this paper still hold. If the possible valuations overlap, however, the line of reasoning developed here does no longer go through. The reason is that if the buyer has a reservation value which is smaller than the highest possible cost of the seller she cannot guarantee herself a strictly positive payoff which is a necessary ingredient of Lemma 4. Furthermore it may no longer be optimal for a low cost seller to mimic the type with the highest possible cost, which is crucial for Proposition 2. Therefore we have to leave the analysis of this case to future research.

APPENDIX

Proof of Lemma 4. Consider any equilibrium and fix a history h_t in which the seller has accepted p_{t+j} iff $p_{t+j} \geq 1$ ($j = 1, \dots, T-t$). Because of $\mu^* > 0$ such a history exists with positive probability. Note that ϵ is independent of T and t and that M has been chosen in a way to guarantee that

¹² See Fudenberg, Levine, and Tirole [3, Lemma 1].

¹³ See Schmidt [14] for a formal proof. A similar result has been shown independently, for example, by Fernandez-Arias and Kofman [1]. They consider the case of a long-run player facing a sequence of short-run players and a uniform distribution of types.

$\varepsilon > 0$. Consider a possibly mixed strategy of the buyer which chooses prices lower than 1 along h_t in at least M periods with positive probability and suppose that the probability of acceptance for each of the first M of these offers is smaller than ε . We will show that this cannot be true in equilibrium, because then the buyer would get less than the lower bound of his equilibrium payoff.

Denote the first M periods in which price offers lower than 1 are offered with positive probability $\tau_1, \tau_2, \dots, \tau_M$. Let $V_{\tau-1}^B(p_\tau, A)$ be the expected equilibrium payoff from the beginning of period $\tau-1$ onwards for the rest of the game, given the history up to but not including period τ and given that in period τ the price p_τ was accepted (R if p_τ was rejected). In period τ_1 , when the first price $p_{\tau_1} < 1$ is offered, the expected payoff for the rest of the game is given by

$$\begin{aligned} V_{\tau_1}^B(p_{\tau_1} < 1) = & \pi_{\tau_1}(p_{\tau_1}) \cdot (b - p_{\tau_1} + \beta \cdot V_{\tau_1-1}^B(p_{\tau_1}, A)) \\ & + (1 - \pi_{\tau_1}(p_{\tau_1})) \cdot \left(0 + \sum_{j=1}^{\tau_1 - \tau_2 - 1} \beta^j \cdot (b - 1) \right) \\ & + (1 - \pi_{\tau_1}(p_{\tau_1})) \cdot \pi_{\tau_2}(p_{\tau_2}) \cdot \beta^{\tau_1 - \tau_2} \cdot (b - p_{\tau_2} + \beta \cdot V_{\tau_2-1}^B(p_{\tau_2}, A)) \\ & + \dots + \\ & + \prod_{j=1}^{M-1} (1 - \pi_{\tau_j}(p_{\tau_j})) \cdot \sum_{j=\tau_1 - \tau_{(M-1)}}^{\tau_1 - \tau_M - 1} \beta^j \cdot (b - 1) \\ & + \prod_{j=1}^{M-1} (1 - \pi_{\tau_j}(p_{\tau_j})) \cdot \pi_{\tau_M}(p_{\tau_M}) \\ & \cdot \beta^{\tau_1 - \tau_M} (b - p_{\tau_M} + \beta \cdot V_{\tau_M-1}^B(p_{\tau_M}, A)) \\ & + \prod_{j=1}^M (1 - \pi_{\tau_j}(p_{\tau_j})) \cdot (0 + \beta^{\tau_1 - \tau_M + 1} \cdot V_{\tau_M-1}^B(p_{\tau_M}, R)). \quad (29) \end{aligned}$$

Note that for all $j, j = 0, 1, \dots, M$,

$$b - p_{\tau_j} + \beta \cdot V_{\tau_j-1}^B(p_{\tau_j}, A) \leq \sum_{j=0}^{\tau_j-1} \beta^j \cdot b < \frac{b}{1-\beta} \quad (30)$$

and that

$$V_{\tau_M-1}^B(p_{\tau_M}, R) \leq \sum_{j=0}^{\tau_M-2} \beta^j \cdot b < \frac{b}{1-\beta}. \quad (31)$$

Now suppose that all $\pi_{\tau_j}(p_{\tau_j}) < \varepsilon, j = 0, 1, \dots, M$. Furthermore we can use that $1 - \pi_\tau(p_\tau) \leq 1$ to get

$$\begin{aligned}
V_{\tau_1}^B(p_{\tau_1} < 1) &< \varepsilon \cdot \frac{b}{1-\beta} + \sum_{j=1}^{\tau_1 - \tau_2 - 1} \beta^j \cdot (b-1) \\
&+ \varepsilon \cdot \beta^{\tau_1 - \tau_2} \cdot \frac{b}{1-\beta} + \dots \\
&+ \sum_{j=\tau_1 - \tau_{(M-1)}}^{\tau_1 - \tau_M - 1} \beta^j \cdot (b-1) \\
&+ \varepsilon \cdot \beta^{\tau_1 - \tau_M} \cdot \frac{b}{1-\beta} + \beta^{\tau_1 - \tau_{M+1}} \cdot \frac{b}{1-\beta}. \quad (32)
\end{aligned}$$

Consider the first and the last term of this expression. It is easy to check that ε has been chosen such that

$$\varepsilon \cdot \frac{b}{1-\beta} + \beta^{\tau_1 - \tau_{M+1}} \cdot \frac{b}{1-\beta} \leq b-1. \quad (33)$$

Given this it is certainly true that $\beta^{\tau_1 - \tau_j} \cdot \varepsilon \cdot (b/(1-\beta)) < \beta^{\tau_1 - \tau_j} \cdot (b-1)$. Therefore we have

$$\begin{aligned}
V_{\tau_1}^B(p_{\tau_1} < 1) &< b-1 + \sum_{j=1}^{\tau_1 - \tau_2 - 1} \beta^j \cdot (b-1) + \beta^{\tau_1 - \tau_2} \cdot (b-1) + \dots \\
&+ \sum_{j=\tau_1 - \tau_{(M-1)}}^{\tau_1 - \tau_M - 1} \beta^j \cdot (b-1) + \beta^{\tau_1 - \tau_M} \cdot (b-1) \\
&= \sum_{j=0}^{\tau_1 - \tau_M} \beta^j \cdot (b-1) \leq \sum_{j=0}^{\tau_1 - 1} \beta^j \cdot (b-1). \quad (34)
\end{aligned}$$

But this is a contradiction to the fact that that we are in equilibrium. By Lemma 3 we know that the right hand side of the inequality is a lower bound for the equilibrium payoff of the buyer. Therefore, if there are M price offers lower than 1, then at least one of them has a probability of acceptance bigger than ε . Q.E.D.

REFERENCES

1. E. FERNANDEZ-ARIAS AND F. KOFMAN, Asymmetric information in repeated bargaining games, mimeo, Sloan School of Management, M.I.T., 1989.
2. E. FUDENBERG AND D. LEVINE, Reputation and equilibrium selection in games with a patient player, *Econometrica* **57** (1989), 759–778.
3. D. FUDENBERG, D. LEVINE, AND J. TIROLE, Infinite horizon models of bargaining with one-sided incomplete information, in "Game Theoretic Models of Bargaining" (A. Roth, Ed.), Cambridge, Cambridge Univ. Press, London/New York, 1985.

4. D. FUDENBERG AND J. TIROLE, Sequential bargaining with incomplete information, *Rev. Econ. Stud.* **50** (1983), 221–247.
5. F. GUL, H. SONNENSCHN, AND R. WILSON, Foundations of dynamic duopoly and the Coase conjecture, *J. Econ. Theory* **39** (1986), 155–190.
6. O. HART AND J. TIROLE, Contract renegotiation and Coasian dynamics, *Rev. Econ. Stud.* **55** (1988), 509–540.
7. D. KREPS, P. MILGROM, J. ROBERTS, AND R. WILSON, Rational cooperation in the finitely repeated prisoner's dilemma, *J. Econ. Theory* **27** (1982), 245–252.
8. D. KREPS AND R. WILSON, Reputation and imperfect information, *J. Econ. Theory* **27** (1982), 253–279.
9. D. KREPS AND R. WILSON, Sequential equilibria, *Econometrica* **50** (1982), 863–894.
10. J. LAFFONT AND J. TIROLE, The dynamics of incentive contracts, *Econometrica* **56** (1988), 1153–1175.
11. E. MASKIN AND J. TIROLE, Markov equilibrium, Discussion Paper, Harvard–M.I.T., 1989.
12. P. MILGROM AND J. ROBERTS, Predation, reputation and entry deterrence, *J. Econ. Theory* **27** (1982), 280–312.
13. K. SCHMIDT, Reputation and equilibrium characterization in repeated games with conflicting interests, *Econometrica* **61** (1993), to appear.
14. K. SCHMIDT, “Commitment in Games with Asymmetric Information,” Ph.D. thesis, University of Bonn, July 1991.
15. J. SOBEL AND I. TAKAHASHI, A multi-stage model of bargaining, *Rev. Econ. Stud.* **50** (1982), 411–426.