Review of Economic Studies (1986), LIII, 709-724 0034-6527/86/00440709$02.00

Non-Cooperative Bargaining Theory: An Introduction

JOHN SUTTON
London School of Economics

The paper provides an informal introduction to some of the main themes of the recent literature on "non-cooperative" or "sequential" bargaining models. It focuses in particular on the relationship between the new approach and the traditional axiomatic approach exemplified by "Nash bargaining theory". It illustrates the new insights offered by the non-cooperative approach, by reference to a detailed analysis of the manner in which the presence of an outside option available to one of the parties will affect the negotiated outcome. Finally, the difficulties which arise in extending this analysis to two-person bargaining with incomplete information, and to n-person bargaining, are discussed.

This is a revised version of the fourth *Review of Economic Studies Lecture* presented in April 1985 at the joint meeting of the Association of University Teachers of Economics and the Royal Economic Society held in Oxford. The choice of lecturer is made by a panel whose members are currently Professors Hahn, Mirrlees and Nobay, and the paper is refereed in the usual way.

1. INTRODUCTION

This paper aims to provide an informal and elementary introduction to an approach to bargaining which has received a great deal of attention over the past few years. The approach involves writing down some particular sequence of moves (offers and replies) to be made over time in the course of negotiations, and then looking for a non-cooperative equilibrium in the game thus specified (in practice, a perfect equilibrium, or in games of incomplete information, a sequential equilibrium).

While the approach appears at first sight to be very different in spirit from the traditional axiomatic approach—in which a bargaining solution is specified by appealing to a number of general requirements which are deemed appropriate on the basis of some *a priori* considerations—the two approaches are in fact complementary. While the complementarity will be illustrated at length in what follows, it may be worth stating the basic point succinctly at the outset: the detailed process of bargaining will differ so widely from one case to another that any useful theory of bargaining must involve some attempt to distil out some simple principles which will hold over a wide range of possible processes. What an axiomatic approach attempts to do is to codify some set of principles of this kind. To design such a set of axioms, though, we need at least to carry out some thought experiments, in order to guide our intuition as to what principles are reasonable, or compelling. The easiest way to do this is to imagine some particular process which might be followed, and to ask whether the principle will hold good in that case. This motivates the idea of looking at some example(s) of non-cooperative games which correspond to a particular process. But we can design many such processes, and the bargaining outcome will depend on the process we choose. Indeed, this is the basis of the assertion which is sometimes made, that the study of such processes is an empty exercise. The alternative view is that the study of such processes may nonetheless allow us to draw out certain simple principles which can guide our (re)formulation of the axioms we use, and inform our application of these axioms. One of the main aims of the present paper is to illustrate this latter viewpoint.

The general idea just stated is not novel; indeed, a concise statement of that view will be found in Nash (1953), in which a particular non-cooperative game (the simultaneous moves "Nash demand game") is employed in order to motivate the Nash bargaining solution. What is novel to the recent literature is the introduction of a new kind of non-cooperative game, which turns out to be particularly helpful in carrying out this programme.

While the first half of the present paper emphasizes the value of this approach, its limitations become more apparent when we turn to the case of incomplete information, and to n-person bargaining. For much of the usefulness of an appeal to non-cooperative models rests on their specifying a unique equilibrium. In the case of two-person bargaining under complete information, an appeal to perfectness is sufficient to ensure uniqueness (Rubinstein (1982)). The imposition of perfectness in this context appears natural, and possibly even compelling. The restrictions which have been employed to ensure uniqueness in games of incomplete information are, however, at best, less persuasive. The case of n-person bargaining games is also problematic. In the second half of the paper, we outline some of the difficulties involved in these areas.

2. RUBINSTEIN'S MODEL

Two agents set out to divide a cake of size 1 between them; if they agree, each receives his agreed share. If they fail to agree, both receive zero.

A simple bargaining process involves the parties taking turns to make proposals: at time 0, player I proposes that he receive some share, x. Player II immediately replies "Yes" or "No". If he says "Yes" the game ends; otherwise, at time 1, Player II makes a proposal to which Player I immediately replies; and so on. The payoff to Player I (Player II) equals his share of the cake as agreed at time t, multiplied by δ_1^t (resp. δ_2^t), where δ_1, δ_2 represent discount factors. To provide some incentive for the players to reach an agreement, we assume $\delta_1, \delta_2 < 1$. A *strategy* for Player I specifies his proposal/reply at each point, as a function of the history of the game up to that point.

It is easy to see that any partition of the cake can be supported as a Nash equilibrium; we proceed to seek a Perfect Equilibrium, viz. a Nash Equilibrium, such that the strategies induced in every subgame form a Nash Equilibrium in that subgame.[1]

It was shown by Rubinstein (1982) that, in this game, there is a unique partition of the cake, which can be supported as a Perfect Equilibrium. In this equilibrium, agreement is immediate, and Player I receives share (payoff) $(1-\delta_2)/(1-\delta_1\delta_2)$, while Player II receives share (payoff) $\delta_2(1-\delta_1)/(1-\delta_1\delta_2)$. To show this, we follow the proof of Shaked and Sutton (1984*a*):

Let M denote the supremum of the share which Player I can obtain in any perfect equilibrium of this game.

TABLE I

Time	Offer made by	I gets at most share	II gets at least share
0	I	$1-\delta_2(1-\delta_1 M)$	
1	II		$1-\delta_1 M$
2	I	M	

Now consider the subgame beginning with an offer made by Player I at time $t=2$. (See Table I.) Notice that this subgame has the same structure as the original game, apart from a rescaling of payoffs, so the supremum of the share which Player I can obtain in any perfect equilibrium of the game is again M. We illustrate this in the bottom row of Table I.

Now consider the offer made by II in the preceding period. Any offer which gives I a share of more than $\delta_1 M$, being the discounted value to him of a share M received one period later, will certainly be accepted. So the share of the cake which II obtains in any perfect equilibrium cannot be less than $1-\delta_1 M$; in fact, this is the infimum of II's share in this subgame.

Now consider I's offer at $t=0$. Any offer by I which gives II a share less than $\delta_2(1-\delta_1 M)$ will certainly be rejected. Hence, I will obtain at most a share of the cake equal to $1-\delta_2(1-\delta_1 M)$; in fact, this represents the supremum of what I will receive in any perfect equilibrium, i.e. it equals M. Setting

$$M=1-\delta_2(1-\delta_1 M),$$

we obtain

$$M=(1-\delta_2)/(1-\delta_1\delta_2).$$

But the preceding argument can now be repeated, but with M defined instead as the infimum of the share received by Player I in any perfect equilibrium of the game; and with the words more/less, most/least, accepted/rejected and supremum/infimum interchanged throughout.

Hence the above equation also defines M as the infimum of the share received by I; thus the shares received in any perfect equilibrium are uniquely defined. Player I receives the share M given by the above equation, and Player II receives $(1-M)$.

It is easy to show that this solution is in fact supported as a perfect equilibrium, and so there exists a unique Perfect Equilibrium Partition. (The strategies are such that a player demands an amount corresponding to this Perfect Equilibrium at each stage, and his rival accepts any demand which does not exceed that amount.) This completes the proof.

A remark is in order regarding the above proof: when we arrive at time $t=2$, we have accumulated a "history" of two moves; the strategies forming a P.E. in the present game might, in principle, specify various P.E.'s in the subgame beginning at $t=2$, as a function of the history of past moves. We do not *assume* here that the strategies supporting the P.E. are stationary (history independent). But by working in terms of the *supremum*, we can exploit the stationarity of the underlying structure.

The above result has the intuitively appealing feature that the more impatient a player, the smaller his share of the cake. It also has the less attractive feature that there is an advantage to "moving first"; for example, where $\delta_1=\delta_2=\delta$, player I receives $1/(1+\delta)$. This advantage can be removed by determining the identity of the proposer, in each period, by "tossing a coin" (Binmore (1980)). An alternative procedure which eliminates the first mover advantage is as follows: let the time delay between successive periods be denoted Δ, and write the discount factor as δ^Δ accordingly. Then in the limit $\Delta\to 0$, the equilibrium share of Player I becomes

$$x=\lim_{\delta\to 0}\frac{1-\delta_2^\Delta}{1-\delta_1^\Delta\delta_2^\Delta}=\frac{\ln\delta_2}{\ln\delta_1+\ln\delta_2}$$

so that $x=\frac{1}{2}$ when $\delta_1=\delta_2$.

An attractive feature of the present approach is that it permits a neat separation between two elements of the bargaining problem: the "institutional structure" within which bargaining occurs (the sequence of moves) and the players' preferences (here, their discount factors). Much of the scope for novel applications afforded by this approach is associated with the judicious design of structures of moves which capture some key feature of the bargaining environment.

Finally, we remark on the sense in which one interprets the "sequence of moves" in this setup. Three views are possible:

(i) We can imagine the delay between successive moves as representing an agent's "speed of response". On this interpretation, a focus on the limit where these delays become arbitrarily short seems compelling.

(ii) We might think of the delay following a proposal by a player as measuring the length of time for which he can commit himself (not to entertain an alternative proposal). Such an approach admittedly begs the question as to what (unmodelled) factors might determine an agent's ability to commit himself. (On the role of commitments, see Crawford (1982).) It also suggests trying to find an analogue for this time delay in actual bargaining institutions.

(iii) It may be best, however, to look on these sequences of moves merely as stylized representations of bargaining processes, within which, by suitably altering the move structure, we may be able to represent certain differences between one bargaining environment and another.

3. AN APPLICATION

One of the main attractions of the type of approach described above is that it permits us to design a wide menu of non-cooperative games (structures of moves), the study of which can in turn help to clarify how certain changes in the bargaining environment might be expected to impinge on the outcome. This point is perhaps best illustrated by pointing to a common confusion which is readily resolved by appeal to a simple extension of the game just considered. The issue concerns the following question: suppose one of the bargaining partners is free to quit bargaining and instead take up some "outside option". How will the value of the outside option available to him impinge on the bargaining outcome? The answer usually given to this is that the presence of an option of value s for Player II "shifts the *status quo* point" from $(0, 0)$ to $(0, s)$, and it is with reference to this new *status quo* point that the Nash bargaining solution should be computed. In our simple setting in which players' utilities are identified with their respective shares of the cake, this reduces to the familiar "split the difference" rule, so that Player II's payoff shifts from $\frac{1}{2}$ to $s+\frac{1}{2}(1-s)$.

Now the non-cooperative game which Nash suggested as a justification for the Nash bargaining solution, the "Nash demand game" (described later), does not appear to correspond well to a situation of the kind just proposed. In what follows, we appeal to some non-cooperative examples in order to draw out a distinction which ought to be made, between those kinds of situation in which "shifting the *status quo* point" seems appropriate, and those in which it does not. Finally, we remark on the relationship between these examples, and the "Nash demand game".

We consider two different games, which are very similar in form, in order to underline as clearly as possible the distinction which is at issue. Both games may be illustrated by reference to the structure of moves shown in Figure 1. (This rather special structure of moves has been designed so as to allow us to obtain our two quite different, but vey simple games, as limiting cases within the same framework.)

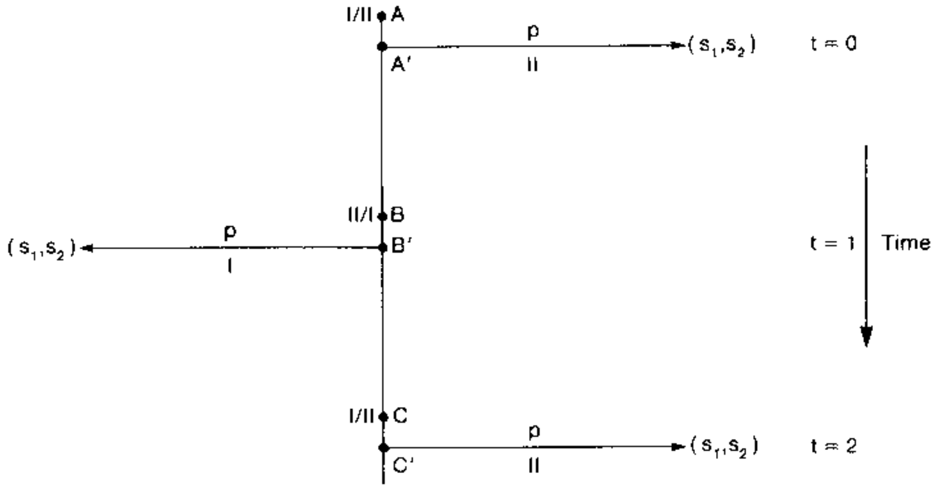

FIGURE 1

Game 1

Let the cake be of size 1, and let both players have a common discount factor $\delta < 1$. At time 0, I makes a proposal to II (node A in Figure 1). II immediately replies Yes or No. If he says Yes, the game ends. Otherwise, a random event occurs with probability p (node A'). If the event occurs, II can choose to quit bargaining, and receive s_2, while I will receive s_1. If the random event fails to occur, the option is not available, and II must await his turn to make a counter-offer (node B). The same pattern of events then recurs as illustrated in Figure 1; following II's offer, at node B, I says Yes or No. In the latter case, he will with probability p be free to quit, whereupon the parties receive discounted payoffs $\delta^t s_1$ and $\delta^t s_2$ respectively. The random events which occur at successive nodes are independent. Payoffs are given by the expected value of the discounted share received.[2] Notice that, if Player I refused to bargain with II and simply waited for his option to become available, he would achieve his security level $S_1 = ps_1 + p(1-p)\delta s_1 + \cdots = ps_1/(1-(1-p)\delta)$. Similarly, II could in the same way attain a payoff of $S_2 = ps_2/(1-(1-p)\delta)$. We suppose the sum $S_1 + S_2$ is less than unity, so that mutually beneficial trade is possible.

The game just specified is readily solved using the method set out above.[3] It has a unique Perfect Equilibrium Partition, whose form depends on the values of s_1, s_2 and p in the following manner:

Define the function

$$\sigma(s_1) = \delta \cdot \frac{1 - ps_i - (1-p)\delta}{1-(1-p)\delta^2}, \qquad i = 1, 2.$$

Then, at equilibrium, agreement is reached at $t = 0$ and Player I receives share x, given by:

$$x = \frac{1}{1+\delta} \qquad s_1, s_2 \leqq \frac{\delta}{1+\delta} \qquad (1.1)$$

$$= \frac{1-\delta+p\delta s_1}{1-(1-p)\delta^2} \qquad s_1 > \frac{\delta}{1+\delta}, \qquad s_2 \leqq \sigma(s_1) \qquad (1.2)$$

$$=\frac{1-(1-p)\delta-ps_2}{1-(1-p)\delta^2} \qquad s_2>\frac{\delta}{1+\delta}, \qquad s_2\leqq\sigma(s_2) \tag{1.3}$$

$$=\frac{1-ps_2-(1-p)\delta+p(1-p)\delta s_1}{1-(1-p)^2\delta^2} \qquad \text{otherwise.} \tag{1.4}$$

This solution takes a relatively simple form in the following limit. Change the time interval between successive offers from 1 to Δ; replace the discount factor δ by δ^{Δ}; finally, replace the probability p by $1-\exp(-\lambda\Delta)$, so as to keep constant the probability per unit time that bargaining is interrupted by the appearance of an outside option. Then, in the limit $\Delta\to 0$, we obtain $\sigma(s_1)\to\frac{1}{2}$ and

$$x=\tfrac{1}{2} \qquad s_1, s_2\leqq\tfrac{1}{2} \tag{2.1}$$

$$=(1-w')\tfrac{1}{2}+w's_1 \qquad s_1\geqq\tfrac{1}{2}, \qquad s_2\leqq\tfrac{1}{2} \tag{2.2}$$

$$=(1-w')\tfrac{1}{2}+w'(1-s_2) \qquad s_2>\tfrac{1}{2}, \qquad s_1\leqq\tfrac{1}{2} \tag{2.3}$$

$$=(1-w)\tfrac{1}{2}+w\{s_1+\tfrac{1}{2}(1-s_1-s_2)\} \qquad \text{otherwise} \tag{2.4}$$

where $w=1/[1-\ln\delta/\lambda\ln\lambda]$, $w'=1/[1-2\ln\delta/\lambda\ln\lambda]$.

What distinguishes these cases is simply that, in the first, where both options are small, Player I strictly prefers to continue bargaining rather than take up his option, when the latter is available; and likewise for Player II. In the second case, Player I prefers to take up his option, rather than continue bargaining; but Player II does not, and so on. In the fourth case, both options are worth taking, when available.

The method of calculation simply involves beginning from node C, and working backwards as before, calculating at each step what the player making an offer needs to propose in order to "buy off" his partner—and noting that this amount depends on whether or not, at each stage, the partner will take his option, if it becomes available.

The weighting factors w and w' ($0\leqq w, w'\leqq 1$) measure the likelihood that the option will be available (i.e. the size of λ relative to the discount factor).

These results are best interpreted by reference to some special cases:

(i) If $s_1=s_2=0$, the game reduces to the Rubinstein game considered above (equations (1.1) and (2.1)).

(ii) Suppose that $s_2>0$, $s_1=0$ and $p=1$. Here, Player II has an outside option of positive value, which is "always" available. This corresponds to the situation posited earlier. From equations (1.1) and (1.3) we have that player I's payoff is

$$x=\min\left\{\frac{\delta}{1+\delta}, 1-s_2\right\}.$$

This corresponds to the case $\lambda\to\infty$, whence $w'\to 1$, in equations (2.1) and (2.3), which then yield $x=\min(\frac{1}{2}, 1-s_2)$.

In other words, either Player II's option exceeds what he would have obtained in the original game, in which Player I needs to offer (marginally more than) s_2 to "buy him off"; or else it does not—in which case the threat of having recourse to the outside option is empty, and it has no effect on the outcome. (For further discussion of this "Outside Option Principle", see Ståhl (1972), Binmore (1986), Shaked and Sutton (1984).)

Game 2

We now turn to an alternative bargaining game, identical to the above, but for one difference. Suppose now that at each node A, B, C, etc., should the random event occur, then bargaining is automatically terminated and the players receive the payoffs s_1 and s_2, suitably discounted.

The solution to this modified game is given by formulae (1.4) and (2.4) above, for *all* values of s_1 and s_2 (i.e. it corresponds to the situation where both options are taken up if available; a moment's reflection makes this obvious).

What is interesting about this latter game is that the pressure on the parties to settle derives not only from their impatience (discount factors), but also from the fear of an exogenously imposed breakdown. The weighting factor w measures the likelihood of such a breakdown relative to the player's discount factor. From (2.4), it follows that as $w \to 0$, we re-establish the "half and half" division of the Rubinstein game. As $w \to 1$, we converge to the "split the difference" rule, $x = s_1 + \frac{1}{2}(1 - s_1 - s_2)$. Hence this rule would be appropriate, for example, in a case where some exogenously imposed intervention, outside the players' control, might cause bargaining to terminate "soon". (This interpretation is due to Binmore, Rubinstein and Wolinsky (1984).)

What matters then is whether access to his outside option affords a player a credible threat, or not. Small options, if chosen voluntarily, have no effect; the "exogenous intervention" mechanism serves to make (even small) threats credible. In the case where both options are taken when offered, we obtain, in both cases, the same result (rearranging (2.4)), i.e. a "generalized" version of the "split-the-difference" rule,

$$x = ws_1 + \tfrac{1}{2}\{1 - ws_1 - ws_2\}. \tag{2.4$'$}$$

Summing up

What we wish to argue here is that an examination of various non-cooperative structures can help us to identify certain simple principles, such as the "Outside Option principle" noted above, which can in turn help in refining the set of axioms to be used. This remains so notwithstanding the fact that results in this area are indeed sensitive to the structure of moves employed. The Outside Option principle is a good example of this, precisely because it follows as a direct consequence of the "perfectness" notion itself: only threats which are credible will have an effect on outcomes. That bargaining agents will in practice fail to be influenced by their opponents' access to some relatively unattractive alternative is of course an empirical issue. Some recent experimental work on the Outside Option principle lends encouraging support to the predictions of the present analysis (Sutton, Shaked and Binmore (1985)).

We complete this section with some informal remarks regarding the relationship between the limiting form of Game 2 above (equation (2.4)), and the Nash Bargaining solution.

A generalization of Game 2 above leads to an interesting connection, which was noted by Binmore (1982), and Binmore, Rubinstein and Wolinsky (1984). Let $\delta = 1$, so that the utility loss derives solely from the risk of breakdown. Let each individual's (hitherto linear) utility function be represented now by the (Von Neuman–Morgenstern) utility functions $u_1(x)$ and $u_2(x)$ respectively, and suppose for simplicity that these functions are smooth and concave. Then it can be shown that the solution to Game 2 above is given by the Nash bargaining solution, i.e. player I receives share

$$x=\operatorname{argmax}\{u_1(x)-u_1(s_1)\}\{u_2(1-x)-u_2(s_2)\}. \tag{3}$$

For the linear case $u_i(x)=x$ this reduces to the "split-the-difference rule" (as obtained in (2.4) with $w=1$).

Given the quite different form of the "Nash demand game" which formed the original non-cooperative underpinning for this solution, it is perhaps worth developing the simple intuition underlying the appearance of this solution point in both games.

Nash's demand game is a modified version of the following simple one-shot game: players I and II simultaneously announce demands x and y, respectively. If $x+y\leqq 1$, each receives the amount demanded; otherwise, both receive zero. Here, any pair (x, y) satisfying $x+y=1$ is a Nash equilibrium. To remove this indeterminacy, Nash modified the game by introducing a "smoothing operation", viz. a function $h(x+y)$ which takes the value 1 on the set $\{x, y \mid x+y\leqq 1\}$, and falls rapidly to zero outside this set. Payoffs are now defined as follows: given demands x and y, player I (resp. II) receives x (resp. y) with probability $h(x+y)$, and receives some amount s_1 (resp. s_2) otherwise. By considering a suitably defined family of functions $h(\cdot)$ which fall off more and more rapidly to zero outside the set $\{x, y \mid x+y\leqq 1\}$, it can be shown that any Nash equilibrium of this game will in the limit coincide with the Nash Bargaining solution (the reader is referred to Binmore (1982) for details).[4] The intuition underlying this result can be seen as follows: given any demand by II, Player I will raise his demand to the point where the marginal utility he obtains from an increment to his share is just offset by the rise in the probability $(1-h)$ of "failure to agree". Now Player i's payoff is:

$$hu_i(x)+(1-h)u_i(s_i)$$

so we must have at equilibrium:

$$(u_i(x)-u_i(s_i))\cdot\frac{dh}{dx}=-hu_i'$$

which immediately implies (3).

Now consider the sequential bargaining story modelled in Game 2 above. We here offer a simple intuitive argument as to why this also leads to the Nash Bargaining solution (3). Refer to Figure 1. When I makes his initial offer at node A, II must, at equilibrium, be indifferent between accepting and rejecting. Player II now balances a small probability p of obtaining "only" s_2, against a small gain from switching roles by having "first move" at B. Let $(x_1, 1-x_1)$ be the equilibrium shares when I has first move,[5] and let $(x_2, 1-x_2)$ be the equilibrium shares when II has first move. Let $z=x_1-x_2$, and suppose the time period to be small, so that the utility increment which I (resp. II) enjoys from moving first is approximately $zu_1'(x_1)$ (resp. $zu_2'(x_2)$). It follows from this that in the limit,

$$p\cdot[u_i(x_i)-u_i(s_i)]\cong zu_i'(x_i), \qquad i=1,2$$

whence, since $x_1\cong x_2$, we again obtain (3).

Thus it is clear why these two apparently very different games, the Nash Demand game, and Game 2 above, lead to the same solution. It is also clear that what "matters" to the result that even small "threats" s_1 and s_2 will affect the outcome, is the fact that both these games build in an enforcement device which makes these threats credible.

All in all, it would seem that a more careful distinction needs to be made, between those situations where "shifting the *status quo* point" is appropriate, and those where it is not, than has been customary in the literature.

4. IMPERFECT INFORMATION

In the "perfect information" models just discussed, the equilibrium strategies imply immediate agreement. Thus a crucial dimension of real bargaining problems is missed. To what extent can the lack of perfect information as to one's partner's preferences, say, lead to equilibrium outcomes involving a period of disagreement?

A considerable literature has developed recently, which rests on a "signalling" argument of the following kind: an agent can identify himself as a "stronger" player—in some appropriate sense—only by making offers and/or replies, which it could not pay a "weaker" player to mimic. Attempts to establish one's "strength" in this way may necessitate the "stronger" player's "holding out" for some time in order to achieve a settlement which, for him, is superior to the early agreement reached by his "weaker" counterpart. Much of the attraction of this kind of argument is that it can be used to produce a rational theory of strike occurrence.

A large number of models of this kind have been investigated in the recent literature, beginning from Fudenberg and Tirole (1983). A full treatment is not attempted here; the present section aims merely to illustrate two difficulties which have attracted particular attention. Both may be seen by reference to the model of Sobel and Takahashi (1983), a model of which is of particular interest, as it forms the basis of a recent econometric study by Fudenberg, Levine and Ruud (1985), in which the model is estimated using the well known Hammermesh-Farber data set on strike duration and wage settlements.[6]

We here confine ourselves to a special case of the Sobel-Takahashi model; the treatment follows Fudenberg, Levine and Ruud (1985).

The model is as follows: a single seller negotiates with a single buyer as to the price of an indivisible good. The seller values the good at zero; this is common knowledge. The buyer imputes some value v, $0 \leqq v \leqq 1$, to the good. The value v is known to the buyer, but not to the seller. The seller's prior belief about v is described by a uniform distribution on $[0, 1]$. In each period $t = 0, 1, 2, \ldots$, the seller announces a price p_t, which the buyer either accepts or rejects. Each party has a discount factor $\delta < 1$; and this is common knowledge.

The model, then, is one of "one sided incomplete information"; only the seller lacks full information. Notice also that, under complete information, the seller, as he makes all the proposals, would receive the entire surplus.

A *system of beliefs* for the seller is a mapping from any history of offers and replies into a probability distribution $F(v|\cdot)$ over the unknown value of v. A *strategy* for the seller is a mapping from (the history of past offers and replies and) his current belief at each time t, into a price p_t. A *strategy* for the buyer is a mapping from (the history of past offers and replies and) the current price offer, into a reply (Yes or No). If a price p_t is accepted in period t, the seller's *payoff* is $\delta^t p_t$ and the buyer's *payoff* is $\delta^t(v - p_t)$.

The seller's belief is updated following each reply made by the buyer, using Bayes' rule. A familiar problem arises here, insofar as some history might unfold, to which the seller had attached a prior probability of zero. It is a simplifying feature of the present model, that this problem can be circumvented, as follows: the only kinds of event to which the seller attaches prior probability zero are:

(i) Following some price offer p_t, the seller expects buyers of all types present to reject p_t; but the buyer instead accepts it. Events of this kind are no problem, for they arise only following an acceptance, in which case the game has ended, and no assignment of posterior beliefs is needed.

(ii) Those for which, following some price offer p_t, the seller expects buyers of all types present to accept p_t. Such an event can only occur if the seller offers a price of zero; and it is readily seen that any strategy involving such a price is *strictly dominated, irrespective of the seller's* beliefs.[7] Hence we may delete such strategies.

We may now define an equilibrium, as a system of beliefs for the seller, and a pair of strategies, such that (i) the seller's beliefs are consistent with Bayes' rule, and (ii) the strategies of buyer and seller are such that following any history of offers (and rejections), then given the seller's updated beliefs, the (induced) strategies employed thereafter by the buyer and seller form a Nash equilibrium. (In the present context this "Bayesian Nash" equilibrium is equivalent to a sequential equilibrium (Kreps and Wilson (1982).)

Suppose the game terminates after T periods. Then, for each T, there exists a unique equilibrium sequence of prices (Sobel and Takahashi (1983)). As $T\to\infty$, this sequence converges to *an*[8] equilibrium sequence of the infinite horizon game, which we now describe.

This equilibrium has the feature that the seller calls a sequence of prices which decline geometrically over time, i.e. $p_t=\gamma^t p_0$, $t=0,1,2,\ldots$, where p_0 and γ are specified below. The buyer will choose to purchase at a time (price) which maximizes $\delta^t(v-p_t)$. (See Figure 2.)

The *buyer's strategy* take the form: accept p_t if and only if $p_t\leqq p^*(v)$, where

$$p^*(v)=\frac{1-\delta}{1-\gamma\delta}v. \tag{4}$$

(Notice that $p^*(v)$ solves the equation $v-p=\delta(v-\gamma p)$.)

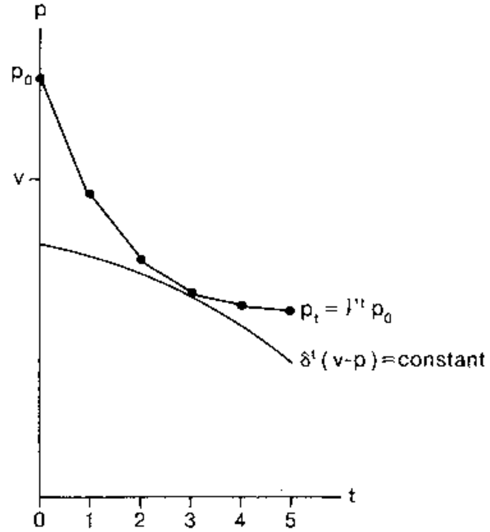

FIGURE 2

The *seller's beliefs* are generated as follows: let $\underline{p}$ be the lowest price offered up to period $t-1$. Then the seller's belief, in any period $t>0$, is described by a uniform distribution on $v \in [0, v^{-1}(\underline{p})]$, where $v^{-1}(\cdot)$ is defined by (4) above. This follows from Bayes' rule, and the buyer's strategy. It further follows that an offer of $p_t \leqq \underline{p}$ in period t will be rejected with probability

$$\frac{v^{-1}(p_t)}{v^{-1}(\underline{p})} = \frac{p_t}{\underline{p}}.$$

So the seller's belief is always described by a uniform distribution on v. As the buyer rejects lower and lower prices, the seller revises downwards his estimate of the highest possible value of v consistent with this behaviour; and the support of the distribution of v shrinks.

To complete our development, we now show that the *seller's strategy* indeed takes the form postulated, and we specify p_0 and γ.

Let $J_t(v)$ be the expected present value of the seller's payoff, when the buyer is known to have a valuation not exceeding v, and let $p_t^*(v)$ be the price for period t specified by the seller's optimal strategy. We must have then, if $p^* \leqq \underline{p}$ (the lowest price offered in the past),

$$J_t(v^{-1}(\underline{p})) = \left(1 - \frac{p_t^*}{\underline{p}}\right) p_t^* + \frac{p_t^*}{\underline{p}} \cdot J_{t+1}(v^{-1}(p_t^*))$$

(while, if $p_t^* > \underline{p}$, we immediately obtain a contradiction). But

$$J_{t+1}(v^{-1}(p_t^*)) = \delta J_t(v^{-1}(p_t^*)) = \delta \frac{p_t^*}{\underline{p}} J_t(v^{-1}(\underline{p}))$$

whence

$$J_t(v^{-1}(\underline{p})) = \frac{\left(1 - \frac{p_t^*}{\underline{p}}\right) p_t^*}{1 - \delta \left(\frac{p_t^*}{\underline{p}}\right)^2}.$$

Choosing p_t^* to maximize the r.h.s., we obtain that $p_t^* = \gamma^t \underline{p}$, where

$$\gamma = (1 - \sqrt{1-\delta})/\delta < 1$$

and so the seller's optimal policy takes the form $p_t = \gamma p_{t-1}$, with $p_0 = \gamma(1-\delta)/(1-\gamma\delta)$.

The model just described captures in a simple manner the intuition suggested earlier.[9] As an explanation of disagreement in bargaining, however, it has a rather worrying feature, as follows: replace the unit time delay between successive offers by Δ. Then, for any fixed time period $\varepsilon > 0$, we can choose Δ sufficiently small, to ensure that agreement will be reached within time ε, with probability $1-\varepsilon$; in other words, agreement is reached "almost immediately".[9] To see this, note that the above solution implies that the fraction of buyers remaining after t rounds is given by $v^{-1}(p_t) = \gamma^{t+1}$, where γ now becomes $(1-\sqrt{1-\delta^\Delta})/\delta^\Delta$, while the number of rounds taking place in time ε equals ε/Δ. Hence the fraction of buyers remaining after time ε equals

$$[1 - \sqrt{1-\delta^\Delta}/\delta^\Delta]^{1+\varepsilon/\Delta}$$

which tends to zero as $\Delta \to 0$. This difficulty was shown to arise by Gul, Sonnenschein and Wilson (1985) not only for the (relatively straightforward) model just considered, but, more importantly, for a (wide class of equilibria of the) modified version of the model in which players take turns to make proposals.[10]

The seriousness of this problem clearly depends on the interpretation one puts on the "structure of moves" employed in these games. On the interpretation that the delay represents the speed of response of agents, the objection appears fundamental. It has been suggested in Section 2 above, however, that there is room for alternative views on this.

Be that as it may, it appears that this feature is intrinsic to models in which only one party lacks complete information. One way of escaping this difficulty may be to go to the case of two sided incomplete information.[11] The extent to which the problem can be circumvented by so doing, however, remains an open question. (For a discussion of this point, see Gul and Sonnenschein (1985).)

Once we turn to games with two sided incomplete information, however, the problem of multiplicity of equilibria becomes much more serious. Even in finite horizon games it will typically be the case that, by changing the beliefs held by agents at those nodes which are off the equilibrium path, a large number of equilibrium outcomes can be obtained. For example, in a two period model, in which each agent can be of two types, Fudenberg and Tirole (1982) fully characterize the set of Bayesian Nash equilibria, showing how alternative assumptions on such beliefs lead to different equilibrium partitions. Rubinstein (1986) provides some striking examples of the delicate manner in which players' beliefs influence outcomes.

The focus of the recent literature in this area has been on a quest for "reasonable" restrictions on beliefs[12] which might lead to a unique outcome. This emphasis is of course a natural one, in the present context; but the degree of success achieved to date has been very limited.

The difficulty stems from the fact that while many plausible restrictions exist,[13] it is not easy to find restrictions which are compelling, but which nonetheless sufficiently narrow the space of equilibria. As an illustration, consider the following criterion, proposed by Kreps, which does indeed seem quite compelling:

Suppose that agent I can be of type a or type b. Then he can convince II that he is of type a by taking an action which satisfies the following two conditions:

(i) if I is of type a, and if his taking this action convinces II that he is of type a, he will be better off;
(ii) if I is of type b, then irrespective of what II is led to believe by his taking this action, then I is worse off.

Grossman and Perry (1984) analyse a game in which players I and II may each be of two types (differing in their respective discount factors). They show that imposing the above criterion does not lead to a unique (sequential equilibrium) partition. They propose a stronger alternative, in which (ii) above becomes:

(ii) if I is of type b, then if II is led to believe by his taking this action that he is of type a, then I is worse off.

This appealing if rather less compelling version of the restriction proves in fact to be "too strong", in the sense that for certain ranges of parameter values (discount factors) there may be no sequential equilibrium which satisfies it. Outside this range of parameter values, however, the equilibrium partition is unique.

The current tendency in this literature, then, is to look for more or less plausible restrictions in order to pin down a unique equilibrium; an approach common to other areas in which games with incomplete information are employed. These problems of non-uniqueness are therefore often seen as reflecting, not a difficulty with the sequential approach to bargaining, *per se*, but rather as a reflection on the general problems inherent in games with incomplete information. The view may well be justified; but consideration of some other difficulties, to which we turn in the next section, suggests that this is not the whole story.

5. *N*-PERSON BARGAINING

The most attractive feature of the Rubinstein model lies in the fact that the natural, perhaps even compelling, restriction of perfectness sufficed to ensure the uniqueness of equilibrium. In view of the general difficulties involved in pinning down solutions in games of incomplete information, it is of particular interest to see whether Rubinstein's uniqueness property survives extensions of the model in other directions.

In the case of n-person bargaining games under complete information, the following counter-example due to Shaked indicates that the perfect equilibrium partition is not in fact unique (the example is discussed fully by Herrero (1985); here, only an outline argument is indicated).

Three players, each with discount factor δ, set out to divide a cake of size 1. Player I proposes a division (a triple (x_1, x_2, x_3) with $x_i \geqq 0$, $\sum x_i = 1$) at time 0: II and III reply in turn; if both II and III agree, the game ends. Otherwise Player II proposes a division at time 2. The game continues in this way, with I, II and III making proposals in successive periods, until agreement is reached.

It can be shown that, for $\delta > \frac{1}{2}$, any partition can be supported as a perfect equilibrium in this game. Say for example we wish to support the partition $(0, 1, 0)$, in which player II receives the entire cake. A set of strategies which supports this can be described informally as follows: player I proposes $(0, 1, 0)$ at $t = 0$. Player II rejects if and only if $x_2 < \delta$. Player III rejects if and only if $\delta \leqq x_2 < 1$. Thus I's initial proposal is accepted at equilibrium. Now if I deviates from his equilibrium proposal, his offer will be vetoed. To ensure that II finds it optimal to veto a proposal of the form $x_2 < \delta$, we arrange that, in the subgame following such a history, II receives the entire cake. To ensure that III finds it optimal to veto if $\delta \leqq x_2 < 1$, we arrange that in the subgame following such a history, III receives the entire cake; and so on (recall that $\delta > \frac{1}{2}$). The details are set out in Figure 3. (The heavy lines in the figure indicate the strategies described here.) At $t = 1$, subgames beginning from nodes labelled B lead to the equilibrium partition $(0, 1, 0)$; those labelled C lead to $(0, 0, 1)$. Two examples are illustrated. At $t = 2$, nodes labelled A' (resp. C') lead to $(1, 0, 0)$ (resp. $(0, 1, 1)$).

The strategies for all subsequent periods are built up following the same principle. Thus II gets the entire cake, because if I offers him less than δ, he will veto it (he then gets payoff δ at $t = 1$), while if I offers him more than δ but less than 1, this proposal will be vetoed by III. In the latter case, the "roles" of the players in the ensuing subgame are switched, I being replaced by II, II by III, and III by I henceforward. Thus III will receive the entire cake at $t = 1$, thus justifying his veto.

A number of comments are in order, concerning this example:

(i) The non-uniqueness does not appear to be removed by any obvious reorganization of the sequence of moves (nor is it removed by requiring replies to be made simultaneously rather than sequentially).
(ii) There is a unique stationary ("history free") equilibrium.

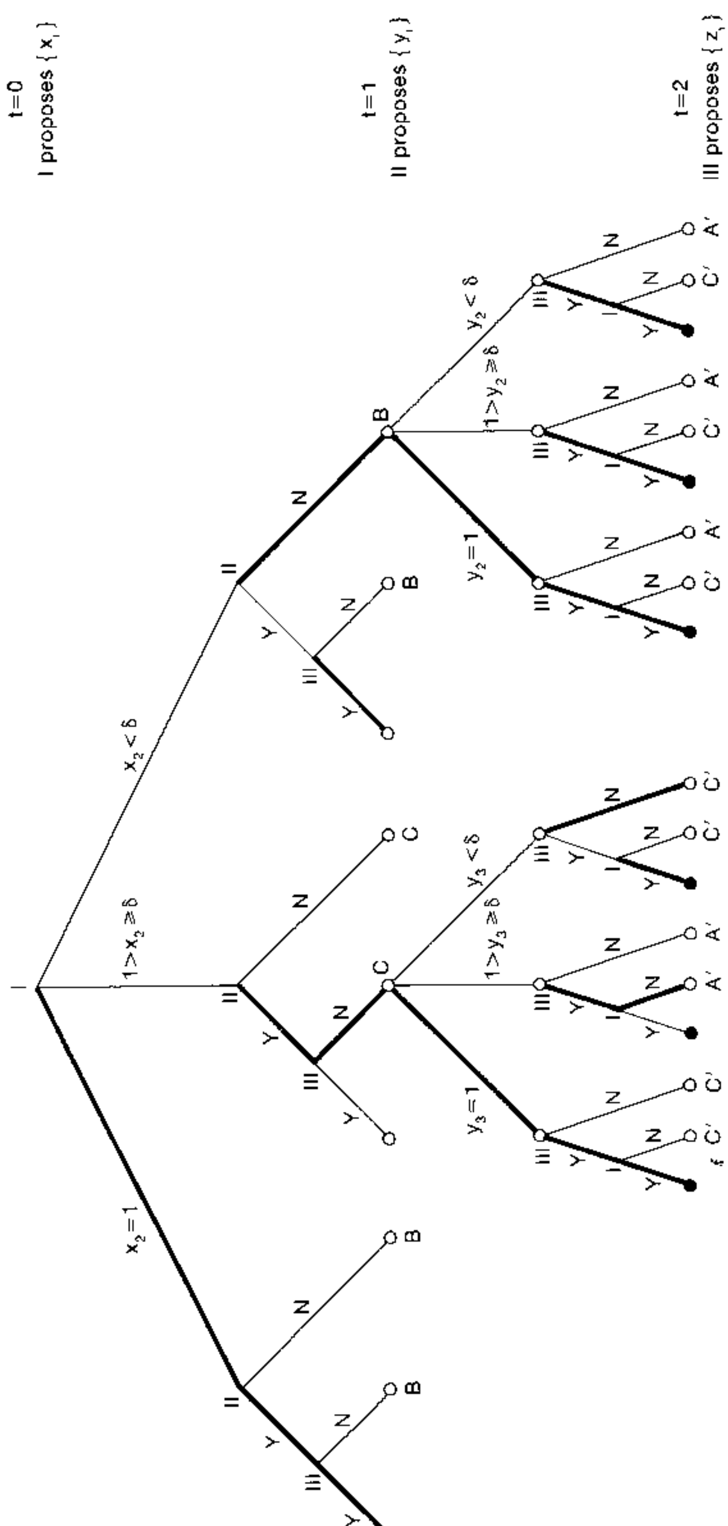

FIGURE 3

The heavy lines show an example of a set of strategies which support an equilibrium in which Player II receives the entire cake

(iii) If the game is truncated after n periods there is a unique perfect equilibrium partition; and the limit of that partition as $n \to \infty$ is that specified by the stationary equilibrium.

(iv) Herrero (1985) demonstrates that the stationary equilibrium is the unique *strong* perfect equilibrium partition.

(v) If strategies are required to vary continuously with the preceding history, then again the only equilibrium is the stationary equilibrium (Binmore (1985)).

It is clear, however, that much remains to be done in this area.

6. CONCLUSIONS

The approach to bargaining which characterizes the papers discussed above appears to be potentially fruitful, in at least two important respects. First, by providing a wide range of non-cooperative models, of a kind which permit us to distinguish salient features of real bargaining problems, it can lead to a more satisfactory mesh between the simple principles embodied in bargaining axioms, and the range of problems which we tackle. Secondly, the incomplete information literature holds out the promise of a richer and more satisfying approach to the analysis of disagreement in bargaining—the problem many economists would see as the central one.

The difficulties encountered in extending the analysis, both to games of incomplete information and to n-person bargaining, are formidable. But if only to keep these difficulties in perspective, it is worth remembering the starting point of the present literature, which lies in Ståhl's early attempts to analyse sequential bargaining processes. It was Ståhl's (1972) examples which prompted Selten's refinement of the Nash equilibrium concept to that of a Perfect Equilibrium, the applications of which far outrun the literature described here. Perhaps the most optimistic view of the present literature is that the sort of difficulties emphasized in the latter half of the present paper may in turn spark off valuable new departures in the analysis of games in extensive form.

First version received November 1985; final version accepted May 1986 (eds.)

In preparing this paper, I have been keenly aware of my considerable debt to my colleagues and collaborators, Ken Binmore and Avner Shaked. I have also benefited from conversations on specific points with Ariel Rubinstein and David Levine, and from the very helpful comments of two referees.

NOTES

1. Strictly what we impose here is subgame perfectness. This suffices to ensure uniqueness. To see that any partition $(x, 1-x)$ can be supported as a Nash Equilibrium, suppose that Player I (resp. Player II) adopts the strategy: always demand share x (resp. $(1-x)$), and accept only if offered at least share x (resp. $(1-x)$).

2. Players are equipped with Von Neuman–Morgenstern utility functions.

3. It can readily be checked that the sup (inf) of I's payoff at A in any P.E.P. is supported by strategies which involve immediate agreement at A. This follows from the assumption $s_1+s_2<1$. Once this point is noted, the proof follows as in Section 2.

4. Nash's original description of h is problematic; more stringent requirements than those proposed in his 1952 paper are needed. For a suitable set of smoothing functions, and a proof that the equilibrium coincides with the Nash Bargaining Solution, see Binmore (1982). In the text, I have for simplicity assumed that h depends on x and y only via their sum. As to whether this non-cooperative game provides a convincing motivation for the Nash Bargaining Solution, see the (critical) comments of Luce and Raiffa (1967, p. 141).

5. Recall that $\delta = 1$, so the game is stationary, even though we have nonlinear utilities.

6. A detailed analysis of the model is developed in Fudenberg, Levine and Tirole (1986), and in Gul, Sonnenschein and Wilson (1985).

7. Let t denote the first time at which a zero price is offered; and let $\underline{p}$ denote the lowest price offered prior to t. All buyers with $v<\underline{p}$ are certainly present at time t. Moreover, some such buyers will strictly prefer to accept any positive price below $(1-\delta)\underline{p}$ at time t, irrespective of the prices prevailing in future periods.

8. It is not the only equilibrium; see Gul, Sonnenschein and Wilson (1985).

9. This "one sided offers" problem is equivalent to that of a durable goods monopolist reducing his price over time in order to "move down his demand schedule" and so price discriminate between consumers who differ in their willingness to pay; the tendency for parties to reach immediate agreement in this bargaining model is directly analogous to the "Coase conjecture" for the durable goods monopolist (Stokey (1982), Bulow (1982)).

10. In the latter model, there may, as noted presently, be many sequential equilibria. The authors show that for *any* such equilibrium which satisfies a mild "Markovian" restriction on beliefs, this property must hold.

11. The examples of equilibria described by Crampton (1984) are suggestive of this possibility.

12. And/or the strategy space. See in particular the analysis of Chaterjee and Samuelson (1984).

13. See for example those of Rubinstein (1985).

REFERENCES

BINMORE, K. G. (1980), "Nash Bargaining Theory I–III", in Binmore, K. G. and Dasgupta, P. (eds.), *Essays in Bargaining Theory* (London: Blackwell) forthcoming.

BINMORE, K. G. (1985), "Modeling Rational Players" (mimeo).

BINMORE, K. G., RUBINSTEIN, A. and WOLINSKY, A. (1985), "The Nash Bargaining Solution in Economic Modelling" (STICERD Discussion Paper No. 112, London School of Economics).

BINMORE, K. G. (1986), "Bargaining and Coalitions", in Roth, A. (ed.), *Game Theoretic Approaches to Bargaining Theory* (Cambridge: Cambridge University Press).

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

CHATERJEE, K. and SAMUELSON, L. (1985), "Bargaining Under Incomplete Information: The Continuum of Offers Case" (mimeo).

CRAMPTON, P. C. (1984), "Bargaining with Incomplete Information: An Infinite-Horizon Model with Two Sided Uncertainty", *Review of Economic Studies*, **51**, 579–593.

CRAWFORD, V. (1982), "A Theory of Disagreement in Bargaining", *Econometrica*, **50**, 607–638.

FUDENBERG, D. and TIROLE, J. (1983), "Sequential Bargaining with Incomplete Information", *Review of Economic Studies*, **50**, 221–247.

FUDENBERG, D., LEVINE, D. and TIROLE, J. (1986), "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information", in Roth, A. (ed.), *Game Theoretic Approaches to Bargaining Theory* (Cambridge: Cambridge University Press).

FUDENBERG, D., LEVINE, D. and RUUD, P. (1985), "Strike Activity and Wage Settlements" (mimeo).

GUL, F. and SONNENSCHEIN, H. (1985), "One Sided Uncertainty Does Not Cause Delay" (mimeo).

GUL, F., SONNENSCHEIN, H. and WILSON, R. (1985), "Foundations of Dynamic Monopoly and the Coase Conjecture" (mimeo).

HARSANYI, J. C. (1967), "Games with Incomplete Information Played by Bayesian Players I–III", *Management Science*, **14**, 159–183, 320–324, 486–502.

HERRERO, M. (1985), "A Strategic Bargaining Approach to Market Institutions" (unpublished Ph.D. thesis, London University).

KREPS, D. and WILSON, R. (1982), "Sequential Equilibria", *Econometrica*, **50**, 863–894.

LUCE, R. D. and RAIFFA, H. (1957), *Games and Decisions: Introduction and Critical Survey* (New York: John Wiley and Son).

NASH, J. (1950), "The Bargaining Problem", *Econometrica*, **18**, 155–162.

NASH, J. (1953), "Two Person Cooperative Games", *Econometrica*, **21**, 128–140.

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

RUBINSTEIN, A. (1985), "A Bargaining Model with Incomplete Information about Time Preferences", *Econometrica*, **53**, 1117–1132.

RUBINSTEIN, A. (1986), "The Choice of Conjectures in a Bargaining Game with Incomplete Information", in Roth, A. (ed.), *Game Theoretic Approaches to Bargaining Theory* (Cambridge: Cambridge University Press).

SELTEN, R. (1985), "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games", *International Journal of Game Theory*, **4**, 25–55.

SHAKED, A. and SUTTON, J. (1984a), "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model", *Econometrica*, **52**, 1351–1364.

SHAKED, A. and SUTTON, J. (1984b), "The Semi Walrasian Economy" (ST/ICERD Working Paper no. 98, London School of Economics).

SOBEL, J. and TAKAHASHI, I. (1983), "A Multi-stage Model of Bargaining", *Review of Economic Studies*, **50**, 411–426.

SUTTON, J., SHAKED, A. and BINMORE, K. (1986), "An Outside Option Experiment", *American Economic Review*, **76**, forthcoming.

STÅHL, I. (1972), *Bargaining Theory* (Stockholm: Economic Research Institute, Stockholm School of Economics).

STOKEY, N. (1982), "Rational Expectations and Durable Goods Pricing", *Bell Journal of Economics*, **12**, 112–128.

LINKED CITATIONS

Non-Cooperative Bargaining Theory: An Introduction
John Sutton
The Review of Economic Studies, Vol. 53, No. 5. (Oct., 1986), pp. 709-724.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198610%2953%3A5%3C709%3ANBTAI%3E2.0.CO%3B2-M

Notes

[9] **Durable-Goods Monopolists**
Jeremy I. Bulow
The Journal of Political Economy, Vol. 90, No. 2. (Apr., 1982), pp. 314-332.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28198204%2990%3A2%3C314%3ADM%3E2.0.CO%3B2-9

References

Durable-Goods Monopolists
Jeremy I. Bulow
The Journal of Political Economy, Vol. 90, No. 2. (Apr., 1982), pp. 314-332.
Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28198204%2990%3A2%3C314%3ADM%3E2.0.CO%3B2-9

A Theory of Disagreement in Bargaining
Vincent P. Crawford
Econometrica, Vol. 50, No. 3. (May, 1982), pp. 607-637.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198205%2950%3A3%3C607%3AATODIB%3E2.0.CO%3B2-Y

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

Sequential Bargaining with Incomplete Information
Drew Fudenberg; Jean Tirole
The Review of Economic Studies, Vol. 50, No. 2. (Apr., 1983), pp. 221-247.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198304%2950%3A2%3C221%3ASBWII%3E2.0.CO%3B2-A

Games with Incomplete Information Played by "Bayesian" Players, I-III. Part I. The Basic Model
John C. Harsanyi
Management Science, Vol. 14, No. 3, Theory Series. (Nov., 1967), pp. 159-182.
Stable URL:
http://links.jstor.org/sici?sici=0025-1909%28196711%2914%3A3%3C159%3AGWIIPB%3E2.0.CO%3B2-P

Sequential Equilibria
David M. Kreps; Robert Wilson
Econometrica, Vol. 50, No. 4. (Jul., 1982), pp. 863-894.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198207%2950%3A4%3C863%3ASE%3E2.0.CO%3B2-4

The Bargaining Problem
John F. Nash, Jr.
Econometrica, Vol. 18, No. 2. (Apr., 1950), pp. 155-162.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28195004%2918%3A2%3C155%3ATBP%3E2.0.CO%3B2-H

Two-Person Cooperative Games
John Nash
Econometrica, Vol. 21, No. 1. (Jan., 1953), pp. 128-140.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28195301%2921%3A1%3C128%3ATCG%3E2.0.CO%3B2-N

Perfect Equilibrium in a Bargaining Model
Ariel Rubinstein
Econometrica, Vol. 50, No. 1. (Jan., 1982), pp. 97-109.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198201%2950%3A1%3C97%3APEIABM%3E2.0.CO%3B2-4

LINKED CITATIONS

Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model
Avner Shaked; John Sutton
Econometrica, Vol. 52, No. 6. (Nov., 1984), pp. 1351-1364.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198411%2952%3A6%3C1351%3AIUAAPE%3E2.0.CO%3B2-4

A Multistage Model of Bargaining
Joel Sobel; Ichiro Takahashi
The Review of Economic Studies, Vol. 50, No. 3. (Jul., 1983), pp. 411-426.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198307%2950%3A3%3C411%3AAMMOB%3E2.0.CO%3B2-D

Vertical Product Differentiation: Some Basic Themes
John Sutton
The American Economic Review, Vol. 76, No. 2, Papers and Proceedings of the Ninety-Eighth Annual Meeting of the American Economic Association. (May, 1986), pp. 393-398.
Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28198605%2976%3A2%3C393%3AVPDSBT%3E2.0.CO%3B2-K

NOTE: *The reference numbering from the original has been maintained in this citation list.*