The Economics of Bargaining

Edited by

Ken Binmore &
Partha Dasgupta

THE ECONOMICS OF BARGAINING

Edited by

Ken Binmore
and Partha Dasgupta

Basil Blackwell

Preface

The Suntory-Toyota International Centre for Economics and Related Disciplines (ICERD) was established at the London School of Economics in 1978. One of its earliest initiatives was the Economic Theory Workshop which was set up as a forum for the interchange of ideas in economic theory and provided with funds from the Suntory–Toyota Foundation to finance international visits and a discussion paper series. Since that time, the Economic Theory Workshop has entertained many visitors, witnessed much heated debate and produced large numbers of discussion papers. A central theme in much of this work has been the theory of games and its applications in economics. In this and the companion volume *Economic Organizations as Games* we offer a selection of papers written by members of the workshops and our visitors on this theme. For each volume, we have written a lengthy, but we hope useful, introduction which is intended to explain the significance of the papers chosen and their place in the general development of the subject.

We should like to express our gratitude to Suntory and Toyota for their vision in supporting this truly international project and to the staff at ICERD for their assistance and guidance in running the workshop. All royalties for this and the companion volume will be donated to ICERD.

Ken Binmore   Partha Dasgupta
London School of Economics   *St John's College, Cambridge*

Acknowledgements

Contents

1

Nash Bargaining Theory: An Introduction

K. Binmore and P. Dasgupta

1 BACKGROUND

In spite of early hopes, it is only in recent years that the von Neumann and Morgenstern theory of games has begun to be genuinely fruitful in economic analysis. In retrospect, it seems clear that the delay was due only in part to the incompleteness of the theory, since much of what is now being used has been available in essence from the early 1950s. A larger stumbling block has been the problem of determining under what circumstances the available theory is (or is not) applicable.

When seeking to employ a game-theoretic analysis, one is faced with two fundamental questions:

  a   What is the appropriate game?
  b   How should it be analysed?

It will typically be impractical to incorporate into the formal model all aspects of the potential interaction between the agents. It is therefore necessary to decide which of these aspects should be formally modelled and which should be taken account of informally. Where and how this line is drawn will be highly significant to the question of what type of analysis is appropriate. If all, or almost all, of the essential strategic structure is adequately formalized, then the conceptual problems involved in determining what constitutes a solution to the game may well be easy. But this will seldom also be true of the technical problems which arise in calculating the solution. It seems doubtful, for example, that it will ever be possible actually to compute the solutions of such games as Chess or 'Go' (although Poker and its relatives are reasonably amenable to analysis). On the other hand, if little of the strategic structure is formalized, then the technical problems are likely to be slight once a solution concept has been chosen. Under the heading of a 'solution concept' it is common to include a very wide variety of ideas. We shall use the term to signify a function which maps from the given formal structure of the game to subsets (or, sometimes, sets of subsets) of possible outcomes. The less

structure incorporated in the formal description of a game, the less complex such functions can be; and therefore the less likely it is that any mathematical difficulties encountered will turn out to be intractable. Equilibrium ideas, such as those discussed in the companion volume *Economic Organizations as Games*, make no sense unless the formal description of the game includes a detailed specification of the strategic opportunities open to the players. Such a specification is often complicated and the corresponding equilibrium analyses are therefore accordingly seldom straightforward. On the other hand, solution concepts like the 'core' or the 'Shapley value' are relatively easy to deal with because their definition requires comparatively little formal structure on the part of a game.

The problem at this end of the spectrum is that it is not clear which of the many available solution concepts, if any, is the 'right' concept for the problem in hand. This question is a difficult one and it is not surprising that authors should be reluctant to provide tightly argued justifications for their answers. But to fudge the question is to risk using game theory as a Procrustean bed on which the economics can be chopped or stretched at will to suit one's computational convenience.[1] This danger is particularly acute in those cases, which frequently occur in a game-theoretic context, where the results of an analysis are surprising or downright counterintuitive. For example, in two-person, pot-limit, straight Poker it is optimal for the opening player *always* to bet on the first round with a ten-high 'bust' (or worse), *never* to bet with a pair of twos and to bet incredibly conservatively with a low straight flush (Cutler 1975). Untutored common sense is therefore not necessarily a reliable guide.

These remarks are not intended to suggest that there is no room in game-theoretic modelling for good judgement or intuitive insight. Quite the contrary is the case. But intuition needs to be trained. In particular, it is important to adopt a critical attitude towards any general assumptions about optimal behaviour which may be advanced and to seek for ways of testing their validity.

The introduction to the companion volume *Economic Organizations as Games*, was insistent on the importance of the idea of a Nash equilibrium. But this idea is only one of several important contributions made by Nash during the short time he devoted to game theory. Perhaps an idea of equal importance is embodied in what has now become known as the 'Nash program'. This provides a conceptual framework within which it is possible to evaluate general assumptions about optimal behaviour in a unified and disciplined manner. The fundamental point is that a choice of solution concept for a game-theoretically based economic model needs to be considered

---

1 In a one-seller, two-buyer market, outcomes in which the buyers form an unbreakable coalition lie outside the core. But this is hardly a reason for excluding them as possible 'economic outcomes' (see Weintraub 1975, p. 70).

very carefully. In so far as bargaining theory is concerned, we shall observe that the appropriate solution concept depends very strongly on the informational and institutional properties of the negotiation arrangements. These, of course, will seldom be deducible from what would usually be regarded as an adequate 'economic' description of a problem. A choice of solution concept therefore implicitly generates assumptions about the 'non-economic' characteristics of the situation the model is intended to describe. It is necessary, at the very least, to be conscious of the fact that such implicit assumptions are being made and to have some feeling for their nature. What is certainly not the case is that the choice of solution concept is simply, or even largely, a matter of taste.[2] If an inappropriate solution concept is grafted onto a model, one has to expect that the results will be grotesquely distorted. The purpose of the Nash program is to minimize the risk of generating such distortions.

2   THE NASH PROGRAM

In the introduction to the companion volume, we introduced the idea of a formal game and commented upon the analysis of such games *in the absence of* a cooperative infrastructure. In this volume, however, our interest centres precisely on those situations in which a cooperative infrastructure *is present.* The necessary considerations are therefore very different.

Firstly, we should clarify what we mean by a 'cooperative infrastructure'. Suppose that it is recognized that a certain situation has game-like characteristics and an attempt is made to construct a formal model G which captures these characteristics. For the purposes of the current discussion it does not matter greatly whether G exhibits all of the properties required of a formal game, because, even if it did, a knowledge of the formal structure of G would still not usually be adequate to analyse the game. In general, it will also be necessary to know something about the *unformalized*[3] activity which may precede the play of the game. It is such unformalized pre-play activity that we have in mind when using the term 'cooperative infrastructure'. (Note the contrast with *contests* which were the focus of attention in the introduction to the companion volume.)

Where possible, for example, players of a game will usually wish to *coordinate* their strategy choices.[4] If some players are reluctant to cooperate perhaps they should be bribed or threatened. If so, how big should the

2 As, for example, is suggested explicitly by Friedman in his otherwise instructive book (1977).

3 Where the opportunities for pre-play activity open to the players are sufficiently rich, it can happen that the precise nature of the *formal* strategies for G becomes almost irrelevant. Descriptions of the formal strategies are then typically omitted.

4 As in the theory of *correlated equilibria* (see Myerson 1984 for a topical account).

bribes be? What threats will be effective? Such pre-play activity is not always permitted. For example, bribery in the species of game played by two contractors submitting tenders to a public official is illegal. Similarly, anti-trust laws are designed to prevent the formation of certain coalitions for mutual advantage. Perhaps of most vital importance is the question of whether any pre-play agreements are *binding* on the players. The interest of the parlour game 'Diplomacy', for example, lies in the judicious breaking of agreements, while, in the commercial world, contracts are supposedly legally enforceable. On a more fundamental level, it may be that the players cannot communicate with each other at all (e.g. if one does not know the other's whereabouts) or that meetings of three or more players are impossible (e.g. if they communicate by telephone).[5] Where self-binding threats are possible, the order in which the players can make their statements becomes highly relevant. This brings us to the general question of commitment. Commitments made physically (for example, by investing in a plant) are easily modelled; but what of verbal commitments? To what extent do these bind those who make them? Can we suppose that commitments the players make are credible because of the impact that a failure to keep the commitment would have on their reputations? If so, what is the mechanism? (For a more detailed discussion of these and related questions on trust, see Dasgupta 1986.)

These considerations are listed in an attempt to make it clear that what happens in a game will in general depend on what may or may not have happened *before* the formal game is played (or on what is anticipated might happen *after* the formal game is played). But solution concepts simply consist of mappings from a formal game G (often suppressing much of the structure of G) to outcomes of G. It follows that the choice of a solution concept must necessarily entail a whole mix of implicit assumptions about the nature of the cooperative infrastructure within which the game is thought to operate. And, to carry the same point further, the existence of a battery of such solution concepts necessarily implies the existence of a *classification* of cooperative infrastructures – each different concept being appropriate for a different infrastructure.

The standard classification, illustrated below, is almost shamefully primitive.
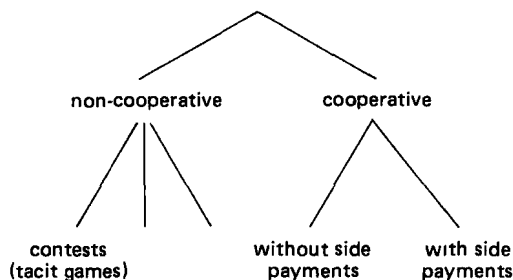
Before strategies are chosen in a formal game played cooperatively, it is supposed that players may communicate costlessly and without restriction and that they may enter into any agreements whatsoever that they choose. Most important of all, it is assumed that mechanisms exist for enforcing such agreements. A side payment is to be thought of as a sum of money which may be paid by one player to another in order to facilitate (or hinder) the signing of a pre-play contract. In principle, the notion of a non-cooperative game covers a wide spectrum of possibilities but attention is usually addressed to what Harsanyi calls 'tacit games' and we prefer to call contests. It is to be understood that, before the play of such a game, the players have no opportunities, explicit or implicit, for any type of communication at all.

This classification is too crude to provide more than preliminary assistance in evaluating solution concepts. For this reason certain concepts for cooperative games come equipped with characterizing 'axiom systems'. Chapter 2 of this volume, for example, provides an alternative axiom system to that usually offered in support of the Nash bargaining solution. But such axiom systems are typically couched in rather abstract terms and there seem to be no cases at all where the appropriate interpretation is not controversial.

The Nash program (outlined in a few sentences by Nash 1951) provides a very much more satisfactory framework within which to consider these matters. His basic view was that the most fundamental type of game is what we have called a *contest*. We have discussed the game-theoretic analysis of such games in the companion volume. Briefly, the notion of a *Nash equilibrium* provides a firm foundation upon which the theory can rest. This is not to say that problems do not exist nor that refinements of the Nash equilibrium idea are not essential.[6] On the contrary, there are numerous difficulties to which we shall find it necessary to return. But these difficulties are of a secondary nature compared with the conceptual problems resolved for contests by the idea of a Nash equilibrium but which remain largely unresolved for other types of game.

We now turn to games other than contests – i.e. games which are to be analysed on the assumption that some measure of pre-play interaction between the players is possible.[7] Nash proposed to deal with the problem of pre-play activity in the following way. If G is a formal game, imagine the various possible steps in the negotiations which precede its play as moves in a larger 'negotiation game' N – i.e. formalize the negotiation procedure. A strategy for the formal game N is a statement of how to conduct the negotiations

---

5 This consideration can matter enormously to market outcomes. (See Binmore 1983, Rubinstein and Wolinsky 1984 and Binmore and Herrero 1985.)

6 We have in mind problems of equilibrium selection and the difficulties which arise in extensive-form games or in games of incomplete information.

7 For simplicity, we leave aside the question of post-play interaction although there will be many situations where this is very significant.

under all possible eventualities and how finally to choose a strategy for G contingent on the course the negotiations took. The negotiation game N is then analysed as a *contest*. The solution of N as a contest then provides a 'solution concept' for the original game G. Nash maintains that *all* 'solution concepts' for G should be regarded this way – i.e. as the solutions of associated negotiation contests.

An obvious objection to this line of attack lies, of course, in the difficulty of finding a model for the negotiation procedure which is both realistic and sufficiently simple to be amenable to analysis. One response is that difficulties cannot successfully be evaded by pretending that they do not exist. However, matters are not so hopeless as they may seem at first sight. In considering appropriate negotiation models, it is necessary to focus on those factors which have, or appear to have, a genuine strategic relevance to the situation. This certainly does not apply to the bulk of manoeuvres common in real-life negotiations. Under this latter heading, for example, come flattery, abuse, the inducement of boredom and other more subtle attempts to put the opponent at a psychological disadvantage. These factors would certainly be of the greatest importance in a behavioural analysis but they have no place in a game-theoretic analysis. A rational player will simply ignore such irrelevancies.[8]

It is as well to enlarge upon this point since it is an issue on which the literature has a tendency to fudge. (See also Section 2 of the introduction to the companion volume.) Consider, for example, a two-person cooperative game. Von Neumann and Morgenstern (1944) originally suggested that the 'solution' of such a game be identified with a whole continuum of possible outcomes called the 'negotiation set' or 'core' of the game. To identify a particular outcome within the core, they suggested, is an enterprise outside the scope of the game theorist on the grounds that the outcome on which the players finally agree will depend on their relative 'bargaining powers' and such *psychological* factors elude adequate formalization. Nash himself makes reference to a similar behavioural notion of 'bargaining skill' in 1950 but explicitly corrects himself in the later paper (1953) where he asserts that, *'with people who are sufficiently intelligent and rational there should not be any question of "bargaining ability", a term which suggests something like skill in duping the other fellow'*. An analogy may be of some help here. Imagine two chimpanzees playing chess. Perhaps in the circumstances, the players might find it difficult to formalize the rules of the game and it will certainly be true that the outcome of the game will depend principally on psychological factors. But these facts are irrelevant to a game theorist who is concerned with what *would* be best for the players if only they knew it.

---

8 Assuming his concern was only with the outcome of G obtained and not with the manner in which it was achieved.

This information will remain the same even if one of our chess-playing chimpanzees is replaced by Capablanca provided only that both have the same preferences over the set of outcomes. A behavioural analysis would be quite a different matter and would doubtless take account of such factors as 'chess playing power' etc. An ethical analysis would similarly require taking into account numerous issues which are irrelevant to a game-theoretic analysis. This is not to say that a behavioural or an ethical analysis might not be more useful than a game-theoretic analysis: only that it is as well not to be confused about what sort of analysis is actually being attempted.[9]

These ideas carry over to bargaining games. If the game-theoretic optimal strategy involves playing like Attila the Hun, then this will remain true no matter who the player might be. It may be true that St Francis of Assisi might find this strategy hard to implement but this is equally true of a chimpanzee attempting the Sicilian Defence. Alternatively, one may object that St Francis would not *want* to behave in this way on the grounds that the 'end does not justify the means'. But this is an ethical question to be dealt with, if at all, by altering the payoffs.[10]

We have been arguing that the Nash program for formalizing negotiation procedures is less formidable than it at first appears provided that one adheres firmly to a game-theoretic viewpoint. Nevertheless, it remains the case that, when the negotiation procedure has been formalized, one must expect to be faced with a structure whose analysis may be technically difficult. For this reason and others, Nash was very far from suggesting that cooperative game theory should be abandoned. To do so would be like attempting to study the movement of billiard balls using only quantum theory. On the contrary, Nash thought of cooperative and non-cooperative theory as complementary and mutually supporting prongs of a pincer attack upon the problem of negotiated games. This view is sometimes classified along with Hegelian dialectic as being impenetrably paradoxical, but we see little justification for such an attitude since the idea, once stated, seems quite simple.

As stressed in Section 1, a mathematical modeller in economic theory (or anything else) has to exercise his judgement in deciding how much of a situation to represent *formally* within his model. His decision will depend on what aspects of the situation are of particular interest to him and on his knowledge of the situation under study. In practice, the extent to which he is successful will depend on how cleverly he draws the line between the formal and the informal so as to maximize the effectiveness of the analytical tools at his

9 The reason for attempting a game-theoretic analysis rather than a behavioural analysis is that it seems sensible to solve the simple problems before seeking to attack the more difficult ones.

10 In a formal game as defined by von Neumann and Morgenstern, the ends are inseparable from the means in that each terminal mode can be reached by only one path through the game tree.

disposal and thus to minimize his dependence on speculation and guesswork. In so far as bargaining is concerned, two immediate issues arise:

(1) The modeller is unlikely to be immediately interested in the details of the negotiations per se. What he will care about is how the *result* of the negotiations is likely to change as relevant economic parameters alter. In many cases he will be seeking only *qualitative* rather than quantitative conclusions on this issue.

(2) The modeller is unlikely to have hard information about a whole spectrum of issues relevant to the detailed conduct of the negotiations. It is remarkable how little of the empirical work on market institutions is directed at issues which a game theorist would regard as significant.

Both these considerations militate in favour of consigning the details of the negotiation process to the area to be dealt with informally by intelligent speculation and informed guesswork. In practical terms, this means choosing an appropriate 'cooperative solution concept' under which heading can be included such notions as Walrasian equilibrium, the Shapley value, the core or a von Neumann and Morgenstern solution set. Within the Nash program, each of these and other cooperative solution concepts are seen as an attempt to describe the outcome of a non-cooperative negotiation game *without formalizing the negotiation procedure.* To be useful, of course, a cooperative solution procedure must be applicable (or thought to be applicable) to a fairly wide class of negotiation procedures. In some cases, such as the Shapley value, the concept is best seen as an 'average' of what is to be expected over such a class of negotiation procedures. In other cases, such as the core or a von Neumann and Morgenstern solution set, the concept does not supply a unique outcome and one is free to attribute this indeterminacy to the breadth of the class of procedures to which the concept is intended to be applicable.

These last remarks make it clear that the Nash viewpoint is valuable as a philosophical prop in thinking about cooperative solution concepts. But the real value of the Nash program lies in the point to be made next. In the physical sciences, experiment provides a final court to which critics of exercises in speculation and guesswork can appeal. In the social sciences, conclusive appeals to experimental evidence are seldom possible – partly because of Murphy's Law[11] and partly because it is seldom feasible to impose the necessary controls. This makes it all the more important in the social sciences to exercise restraint in making assumptions – particularly those which are not stated explicitly in the description of the formal model but appear implicitly as a result of the choice of a structure for the formal model and whose justification, if any, is based on informal considerations. One form of self-discipline available to a mathematical modeller in both the physical and the social

11 Under carefully controlled experimental conditions, laboratory animals do what they damn well please.

sciences is the mind experiment[12] and it is the idea of a mind experiment which is at the heart of the non-cooperative side of the Nash program.

The purpose of constructing negotiation contests in which the details of the negotiation procedure are spelled out explicitly and without omission is *not* because it is thought that such models will displace the use of cooperative solution concepts. The purpose is to *test* cooperative solution concepts.[13] If a certain concept is claimed to apply in the presence of a particular type of negotiation set-up, one may ask the claimant for his reasons. Typically, the reasons offered will constitute a defence not only for the negotiation set-up originally contemplated but also for other related procedures. The argument can then be examined by formalizing the simplest of these procedures and analysing the resulting contest non-cooperatively. Such a method allows the exposure of silly theories (of which there seems no lack) and the opportunity to refine incomplete theories.[14] *The hallmark and the essence of the Nash program is therefore the imperative to test abstract or informal reasoning with simple but specific negotiation models.* And an important duty for games theorists is to provide a battery of suitable models to facilitate such testing.

As an example, consider the use that De Menil (1971) and others have made of the Nash bargaining solution in discussing union power and related issues. From the point of view of the Nash program, such work is suspect because it does not seriously address the question of *why* the Nash bargaining solution is to be used rather than some other cooperative notion. Like ready-made suits, solution concepts should not be bought without first trying them on for size. As it happens, recent work, based on an important paper by Rubinstein (chapter 3 in the current volume), indicates that such suspicion is justified. The difficulty concerns 'outside options' – by which we mean the utility levels which bargainers can achieve by leaving the bargaining table and abandoning the negotiations. Traditional wisdom takes for granted that these should be dealt with by identifying the 'status quo' point with the 'outside option' point. However, an examination of the appropriate extension of Rubinstein's very natural, non-cooperative bargaining model shows that, while this identification may have merit in some special circumstances, it is *not* the right way to proceed in general.[15] The point is that

12 This point is made at greater length in the Introduction to the companion volume, *Economic Organizations as Games*.

13 Or, on occasion, to assist in the formulation of new cooperative solution concepts.

14 By the type of process, long familiar to mathematicians, described by Lakatos in his *Proofs and Refutations* (1976).

15 See Binmore (1983) or Shaked and Sutton (1984). A more widely applicable procedure is to locate the 'status quo' at the point which would result if nobody left the table but there was perpetual disagreement and to observe that 'outside options' will affect the bargaining outcome if and only if it assigns one or more of the bargainers less than his outside option utility level.

traditional wisdom, which is intuitively quite compelling on this question, would not have been challenged, let alone altered, without the intervention of a testing mind experiment as advocated by the Nash program.

3   THE FORMALIZATION OF NEGOTIATION PROCESSES

At first sight, it is surprising how difficult it can be to identify the crucial features of a bargaining process. One would think that the very considerable experience we all have of bargaining at one level or another would be more useful in this respect than it seems to be.[16] This apparent failure of our intuitive insight makes the availability of the von Neumann and Morgenstern formulation of an extensive game particularly valuable. The attempt to fit a negotiation process into their framework forces attention to be paid to features of the process which otherwise might be dismissed as irrelevant. One of the more significant discoveries of modern bargaining theory has been the identification of *discontinuities* in the bargaining outcome as certain 'frictions' or 'imperfections' in the bargaining process are allowed to tend to zero. In the limit, a unique outcome is obtained in many cases. On the other hand, if the frictions are set equal to zero a priori, then a whole range of outcomes appears possible. What seemed at the time to be an entirely natural and straightforward idealization therefore now appears as a possible root cause for the indeterminacy problem which has plagued bargaining theory ever since Edgeworth wrote on the subject. (See chapters 4, 5 and 8.)

There is still much to be learned about bargaining processes but it is now fairly well-established that one cannot usually offer a sensible estimate of what is likely to happen without having a view on the roles to be ascribed to

    a   commitment;
    b   time;
    c   information

within the bargaining process. We shall take these one at a time.

16 It is easy, if not particularly profitable, to speculate on the reasons. The simplest explanation is that psychological issues dominate at the expense of rational issues (as in poker games among novices). Possibly, social evolution has not equipped us well for dealing with bargaining situations in isolation since we shall usually have to deal with our bargaining partners in the future on a variety of different issues. Possibly the problem is simply that we are reluctant to think too closely about our bargaining behaviour in case we find a second-hand car-dealer lurking inside our skulls. Where personal issues are concerned, we certainly have a strong tendency to pretend that we have not been bargaining at all. A more encouraging hypothesis for a game theorist is that our bargaining behaviour is shaped by what *would happen* if we did not bargain as we do. Such an explanation requires taking account of events which are not normally observed and are unlikely to be properly understood.

3.1 Commitment

In much economic theory it is taken for granted that agents can write whatever contracts they choose and that these contracts will be totally binding on the players – i.e. that the basic hypothesis for a cooperative game is satisfied. Two defences of the hypothesis are usually advanced. The first cites the existence of a legal system and the second refers to the importance of maintaining a reputation for good faith in so far as future transactions are concerned. In many cases both considerations will be relevant. Whatever the reasons, all the essays in this volume take for granted that the aim of bargaining is to provide a contract describing the agreed transaction and that all agents regard such a signed contract as equivalent to the transaction itself. Of course, there will be situations for which this is a bad assumption[17] as the current disarray in the world banking system amply demonstrates.[18] However, the hypothesis that the final contract reached will be honoured does not seem an unreasonable assumption from which to begin.

But it does not follow that, because it is reasonable to assume that one type of contract will be honoured, then *any* type of contract will be honoured. In particular, it is easy to overestimate what can be done via the legal system in this context. We shall take up the issue of contracts which relate to the actual transaction again under the heading of 'information'. For the moment, however, we wish to concentrate on commitments (unilateral or multilateral) that the players may wish to make about how they propose to conduct the negotiations in the future. Schelling (1960) has emphasized both how useful it is to players if they can leave their opponent with a genuine 'take it or leave it' problem and, simultaneously, how hard it is in practical terms to make the necessary commitments stick. Perhaps the best known example, although in a somewhat different context, is that of Ellsberg (1975). He cites the kidnapping victim who would dearly wish to make a commitment not to reveal the identity of the kidnapper in order to escape murder but has no way to do so.

Of course, it does not follow that, because the legal system is unlikely to be effective in enforcing in-bargaining commitments, that such commitments may not be enforced because of the damage that a failure to carry through on a commitment may incur in respect of a player's reputation. Obviously, a reputation for 'toughness' can be very valuable if bargaining with an individual with a reputation for 'softness'. This is another matter we shall take up again under the heading of 'information'. For the moment,

17 One could, of course, regard the players as bargaining over their valuations of the possible contracts or restrict attention to 'self-enforcing' contracts in which goods are exchanged a little at a time (as in illicit drug dealing).

18 Or, less obviously, the Westinghouse contracts on uranium supply.