D S E

Dipartimento   **Scienze**   **Economiche**

Working Paper

Department of Economics

Ca' Foscari University of Venice

Carlo Carraro

Carmen Marchiori

Alessandra Sgobbi

Advances in negotiation theory: bargaining, coalitions and fairness

Advances in Negotiation Theory:
Bargaining, Coalitions and Fairness

Carlo Carraro

University of Venice, Fondazione Eni Enrico Mattei, CEPR, CESifo and CEPS

Carmen Marchiori

London School of Economics and Fondazione Eni Enrico Mattei

Alessandra Sgobbi

SSAV and Fondazione Eni Enrico Mattei

Abstract

Bargaining is ubiquitous in real-life. It is a major dimension of political and business activities. It appears at the international level, when governments negotiate on matters ranging from economic issues (such as the removal of trade barriers), to global security (such as fighting against terrorism) to environmental and related issues (e.g. climate change control). What factors determine the outcome of negotiations such as those mentioned above? What strategies can help reach an agreement? How should the parties involved divide the gains from cooperation? With whom will one make alliances? This paper addresses these questions by focusing on a non-cooperative approach to negotiations, which is particularly relevant for the study of international negotiations. By reviewing non-cooperative bargaining theory, non-cooperative coalition theory, and the theory of fair division, this paper will try to identify the connection among these different facets of the same problem in an attempt to facilitate the progress towards a unified framework.

Address for correspondence:
Carlo Carraro
Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S.Giobbe
30121 Venezia - Italy
Phone: (++39) 041 2349166
Fax: (++39) 041 2349176
e-mail: ccarraro@unive.it

Table of Contents

0

1 Introduction

Bargaining is ubiquitous in real-life. In the arena of social interaction, for example, a married couple is almost constantly involved in negotiation processes throughout the relationship, from the decision of who will look after the children, to the question of whether or not to buy a house, how to manage the resources of the family and so on. In the political arena, a bargaining situation exists, for example, when no single political party on its own can form a government, but different parties have to make alliances and agree on a common program for them to have the chance of winning. At an international level, governments are often engaged in a variety of negotiations on matters ranging from economic issues (such as the removal of trade barriers), to global security (such as fighting against terrorism) to environmental and related issues (such as pollutant's emission reduction, water resource management, biodiversity conservation, climate change control, etc.).

What factors determine the outcome of negotiations such as those mentioned above? What strategies can help reach an agreement? How should the parties involved divide the gains from cooperation? With whom will one make alliances?

The study of any bargaining process is extremely hard, involving a multiplicity of questions and complex issues. As a consequence, the research literature in this field has not yet been able to develop a comprehensive framework for analysis, and a number of theories have been proposed instead, each focusing on single aspects of the problem.

So, for instance, the issue of how to divide the payoffs from cooperation among the parties is traditionally addressed within cooperative bargaining theory, which makes, in turn, "beneficial" assumptions about which properties the equilibrium allocation should have, and does not explicitly address the question of which strategies will be adopted by the negotiators.

In many real-life situations, however, cooperation cannot be ensured, and binding agreements are not a feasible option. Therefore, the strategic choices of the actors involved in the bargaining process need to be explicitly modelled in order to determine the final outcome of the negotiation. *Non-cooperative bargaining theory* is more concerned with these situations and focuses on the bargaining procedures in the attempt to determine which equilibrium outcome will prevail in the absence of interventions.

1

When multiple players are involved in the bargaining, there is the possibility that coalitions form. Traditional bargaining theory is not suitable for representing such situation because it is based on the assumption that only two possible outcomes can arise: the fully cooperative outcome and the fully non-cooperative outcome, where respectively an agreement among all parties is reached and no agreement forms. *Non-cooperative coalition theory* considers this interesting aspect of negotiation processes and, without making any assumption on the final result, analyses the incentives which players may have to form coalitions, and how they may affect the final outcome of the negotiation. The study of coalition formation is particularly important in bargaining contexts where positive externalities are present. In this case, due to players' incentive to free ride, it is quite unlikely that the 'grand coalition' will form and 'partial agreements' usually arise.

Finally, traditional models of negotiation have focused almost exclusively on the efficiency properties of both the process and the outcomes. Yet, as every day experience indicates, considerations other than efficiency play a crucial role in selecting which agreement will be reached – if any at all – and through which path. The *theory of fair division* focuses on processes and strategies which respond not only to Pareto efficiency, but also to equity, envy-freeness, and invulnerability to strategic manipulation.

Whilst the theoretical literature offers this classification into different approaches to the same problem, in the applications the division is not so clear-cut. The lack of a unified theoretical framework to address negotiations has meant that the various, isolated, parts of the theory have been of little empirical use.

Whilst recognizing the importance of cooperative game theory, this paper will mainly focus on a non-cooperative approach to negotiations, which is particularly relevant for the study of international negotiations. In particular, by reviewing non-cooperative bargaining theory, non-cooperative coalition theory and the theory of fair division, this survey will try to identify the connection among these different facets of the same problem in an attempt to facilitate the progress towards a unified framework.

The structure of the paper is as follows. Section 2 will briefly describe the principles of cooperative bargaining theory, which represent the origins of formal theory of bargaining, and will discuss the links between the cooperative and non-cooperative approaches in order to introduce and motivate the study of non-

cooperative bargaining. Section 3 will analyse in detail the famous alternating-offer game proposed by Rubinstein in 1982, which constitutes the starting point for non-cooperative bargaining theory. Section 4 will then discuss some important extensions of this model. Section 5 will be devoted to non-cooperative coalition theory, with the intent of providing insights into its latest developments, which seek to link the theory of coalition formation to the bargaining process. Section 6 will be concerned with the theory of fair division. In particular, the questions will be addressed of how fairness considerations can alter the results of the standard theoretical models, and how fair division algorithms can be incorporated in the existing theories. Finally, some concluding comments will be provided in Section 7.

2   Cooperative versus Non-cooperative Bargaining theory

The formal theory of bargaining originated in the early 1950s with John Nash's work, which establishes the basic framework of the 'axiomatic (or cooperative) approach' to negotiations.

Following Nash, a 'bargaining situation' can be defined as a situation in which *(i)* individuals (or "players") have the possibility of concluding a mutually beneficial agreement, *(ii)* there is a conflict of interests about which agreement to conclude, and *(iii)* no agreement may be imposed on any individual without his approval.

More precisely, Nash defines a *'bargaining problem'* to be the set of utility pairs that can be derived from possible agreements, together with a pair of utilities which is designated to be the 'disagreement point'. This idea can be exemplified with the help of a diagram. Figure 1 depicts a bargaining situation in which two players (*A* and *B*), whose utilities are measured along the *x* and *y*-axis respectively, bargain over the partition of a single cake of known size. The point denoted (d_A, d_B) is the point of no agreement – and it determines the minimum level of utility each party is ready to accept. All points to the North East of (d_A, d_B) represent an improvement for both players and, together, they define the negotiation set[1]. No agreement above the frontier is feasible, and all points on the frontier to the North East of the no-agreement point are Pareto efficient (that is, no player can be made

---

[1] Or zone of agreement, or bargaining set.

better off without the other player being made worse off by moving away from such a point).

Figure 1: The bargaining problem – zone of agreement conceptualisation

Utility of player B

Negotiation set

Optimal Pareto set

(d_A, d_B)

Utility of player A
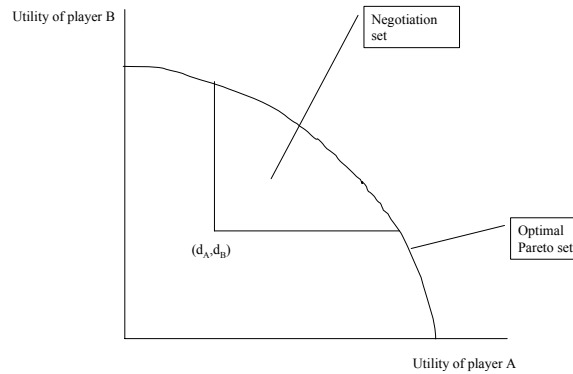
A *'bargaining solution'* is a function (or formula) that assigns a single outcome to every such problem. The Nash bargaining solution is derived from a number of *axioms* about the properties that it would seem natural for the negotiation outcome to have.

In particular, Nash proposes that a bargaining solution should satisfy the following four axioms:

Ax1: Scale Invariance, that is, monotone transformation of the utility functions should not alter the bargaining solution.

Ax2: Symmetry: players are identical, and so they are interchangeable. All differences should be taken care of in the definition of the bargaining set and disagreement points.

Ax3: Independence of irrelevant alternatives, that is, the exclusion of non-selected alternatives from the bargaining set should not alter the bargaining solution.

Ax4: Pareto efficiency: the solution should be Pareto efficient.

It turns out that there is precisely one bargaining solution satisfying these four axioms, and this solution has a very simple functional form: it selects the

utility pair that maximises the product of the players' gains in utility over the disagreement outcome.

$$\max_{u_A, u_B} \left(u_A - d_A\right)\left(u_B - d_B\right)$$

Having axiomatically identified the equilibrium solution, cooperative bargaining theory then concentrates on the problem of how to divide the benefits from agreement among the negotiating parties. [2]

A limit of this approach is that it does not capture the details of the bargaining process. In other words, the process required to arrive at the final outcome is left un-modelled. The justification for this is that rational actors will always choose the outcome that maximises their value. The most efficient solution, therefore, will always be realised regardless of the process.

In fact, as pointed out in the introduction, in many real-life situations cooperation cannot be ensured and binding agreements are not a feasible option because of the absence of a legitimate authority which can impose a centralised solution and/or the complexity of the bargaining situation often involving many parties with very different interests. In such contexts, the strategic choices of the actors involved in the bargaining process need to be explicitly modelled in order to determine the final outcome of the negotiation.

Non-cooperative bargaining theory, which is the focus of this review, is more concerned with these situations and analyses exactly the bargaining procedures, in the attempt to find theoretical predictions of what agreement, if any, will be reached by the bargainers. In particular, this approach seeks to identify the strategies that may sustain cooperation and the variables that may influence agents' behaviour, such as bargaining power, incomplete information, and power relations. In the next two sections, the fundamentals of non-cooperative bargaining theory and its main extensions will be analysed and discussed in details.

---

[2] The research literature on cooperative bargaining theory have proposed different approaches and solutions to the basic problem analysed by Nash. Several extensions have also been developed. See Patrone *et al.* (2004) for a recent survey.

3   Fundamentals of Non-cooperative Bargaining Theory. The basic Rubinstein Alternating-Offer Game

The seminal paper by Rubinstein (1982) represents the starting point for formal non-cooperative bargaining theory. The model developed in this work proposes an attractive and intuitive process of bargaining, and provides a basic framework which can be adapted to many economic and non-economic situations. Sections 3.1 to 3.3 will present and discuss the general structure of the game and its main results, while section 4 will analyse some important extensions of the model.

3.1   Structure of the game

The situation modelled by Rubinstein is the following. There are _two_ players $i = 1,2$ who bargain over a _single_ 'pie' of size 1. An agreement is defined as a pair (x_1, x_2), where x_i is Player i's share of the pie, and the set of possible agreement is:

$$X = \left\{ (x_1, x_2) \in R^2 : x_1 + x_2 = 1 \text{ and } x_i \geq 0 \text{ for } i = 1,2 \right\}$$

The players' preferences over X are diametrically opposed. Each player is concerned only about the share of the pie that he receives, and prefers to receive more rather than less. That is, for $i=1,2$ player i prefers $x = (x_1, x_2) \in X$ to $y = (y_1, y_2) \in X$ if and only if $x_i > y_i$.

The bargaining procedure is as follows. Players can take actions only at times in the (infinite) set $T = \{0, 1, 2, ...\}$. In each period $t \in T$ one of the players, say _1,_ proposes an agreement and the other player (_2_) either accepts the offer or rejects it. If the offer is accepted, then the bargaining ends, and the agreement is implemented. If the proposal is rejected, then the play passes to period (_t+1_), where Player _2_ proposes an agreement and Player _1_ in turn accepts or rejects. The game continues in this way indefinitely until an offer is accepted. At all times, each player knows all his previous moves and all those of her opponent, then a _complete information_ scenario is assumed.

The first two periods of the game are shown in

Figure 2. Play begins at the top of the tree, and time starts at period 0. Player 1 is the first to move and he has a continuum of choices which corresponds to the agreements (members of X) he can propose. Each possible proposal leads to

a decision node for Player 2, at which she accepts (A) or rejects (R) the offer. If Player 2 agrees (right-hand branch), then the game ends and the agreement $x=(x_1, x_2)$ is reached at time $t=0$. If Player 2 rejects the offer (left-hand branch), then play passes to period 1, when it is Player 2's turn to make a proposal. A typical offer of Player 2 is $y=(y_1, y_2)$; for each such offer Player 1 says A(ccept) or R(eject). If he chooses A, the game ends with the outcome y at $t=1$; if he chooses R, then the game continues, Player 1 makes a further offer, Player 2 responds, and so on.

Figure 2: The first two periods of the basic Rubinstein's alternating-offer game

3.2 Assumptions on players' preferences

To complete the description of the game, we need to specify a number of assumptions. Rubinstein assumes that each player $i=1,2$ has a complete transitive reflexive preference ordering ϕ_i over the set $(X \text{ x } T) \cup \{D\}$[3] of *outcomes* and that the players' preference orderings satisfy the following conditions:

---

[3] *D* represents the disagreement point.

A1. (*Disagreement is the worst outcome*). For every $(x = (x_1, x_2), t) \in X \times T$ and $i=1,2$, we have $(x = (x_1, x_2), t) \geq_i D = (d_1, d_2)$.

A2. (*Pie is desirable*). For any $t \in T, x \in X, y \in X$ and $i=1,2$ we have $(x = (x_1, x_2), t) \phi_i (y = (y_1, y_2), t)$ if and only if $x_i > y_i$.

A3. (*Time is valuable*). For any $t \in T, s \in T, x \in X$ and $i=1,2$ we have $(x = (x_1, x_2), t) \phi_i (x = (x_1, x_2), s)$ if $t < s$ (and $x_i > 0$).

A4. (*Continuity*). Player *i*'s preference ordering is continuous.

A.5. (*Stationarity*). For any $t \in T, x \in X, y \in X$ and $i=1,2$ we have $(x = (x_1, x_2), t) \phi_i (y = (y_1, y_2), t+1)$ if and only if $(x = (x_1, x_2), 0) \phi_i (y = (y_1, y_2), 1)$.

A6. (*Increasing loss to delay*). The difference $v_i(x_i)|_{t=1}$, where $v_i(x_i)|_{t=1}$ is the 'present value' of $(x_i, t=1)$ for player *i*, is an increasing function of x_i.

The first assumption concerns the 'disagreement point', *D*, and requires that this is the least-preferred outcome for both players. The remaining conditions concern the behaviour of preferences on the space *X* x *T*. First of all, it is required that among agreements reached in the same period, Player *i* prefers larger value of x_i (*A2*) and prefers to obtain any given partition of the cake sooner rather than later (*A3*). Assumption *A5* is then introduced in order to simplify the structure of preferences. It requires, indeed, that the preferences between $(x=(x_1,x_2), t)$ and $(y=(y_1,y_2), s)$ depend only on *x, y,* and the *difference s – t*. The final condition, *A6*, states that the loss to delay associated with any given amount is an increasing function of the amount.

An example of utility function for which conditions *A1* through *A6* are satisfied is the following

$U_i(x_i, t) = \delta_i^t x_i$, where $\delta_i \in (0,1)$ is player *i*'s discount factor[4]. The preferences represented by this utility function are traditionally called *time preferences with a constant discount rate*.

---

[4] More precisely, we have $\delta_i = exp(-r_i t\varDelta)$, where r_i is player *i*'s discount rate and $r_i > 0$. Therefore, if the discount rate *r* decreases, then the discount factor *δ* increases. This means that player *i* cares more about the future and therefore becomes more patient.

3.3 Main Results

The Equilibrium of the game

Rubinstein (1982) proves that *every bargaining game of alternating offers in which players' preferences satisfy A1 through A6 has a unique subgame perfect equilibrium (SPE)*[5]. In correspondence to this equilibrium:

- Player 1 proposes the agreement $x^* = (x_1^*, x_2^*)$, defined in equation (2.1) below, whenever it is his turn to make an offer, and accepts an offer $y = (y_1, y_2)$ of Player 2 if and only if $y_1 \geq y_1^*$;
- Player 2 always proposes $y^* = (y_1^*, y_2^*)$, whenever it is her turn to make an offer, and accepts an offer $x = (x_1, x_2)$ of Player 1 if and only if $x_2 \geq x_2^*$.

The outcome is that Player 1 proposes $x^* = (x_1^*, x_2^*)$ in period 0, and Player 2 immediately accepts this offer.

In particular, the SPE of the game corresponds to the unique solution of the following equations:

$$y_1^* = v_1(x_1^*)|_{t=1} \qquad \text{and} \qquad x_2^* = v_2(y_2^*)|_{t=1}$$
(2.1)

where the functions $v_1(x_1^*)|_{t=1}$ and $v_2(y_2^*)|_{t=1}$ represent respectively the present value of $(x_1^*, t=1)$ for Player 1 and the present value of $(y_2^*, t=1)$ for Player 2.

In the case of time preferences with constant discount rates (i.e. Player i's preferences over outcomes $(x = (x_1, x_2), t)$ are represented by the utility function $U_i(x_i, t) = \delta_i^t x_i$), (2.1) implies that $y_1^* = \delta_1 x_1^*$ and $x_2^* = \delta_2 y_2^*$, so that

$$x^* = \left(\frac{1 - \delta_2}{1 - \delta_1 \delta_2}, \frac{\delta_2(1 - \delta_1)}{1 - \delta_1 \delta_2} \right) \qquad \text{and} \qquad y^* = \left(\frac{\delta_1(1 - \delta_2)}{1 - \delta_1 \delta_2}, \frac{1 - \delta_1}{1 - \delta_1 \delta_2} \right)$$

(2.2)

Thus, if $\delta_1 = \delta_2 = \delta$ (that is, if the discount factors are equal), then $x^* = (x_1^*, x_2^*) = \left(\frac{1}{1 + \delta}, \frac{\delta}{1 + \delta} \right)$.

---

[5] A strategy pair is a *subgame perfect equilibrium (SPE)* of a bargaining game of alternating offers if the strategy pair it induces in every subgame is a Nash equilibrium of that subgame.

It is important to notice that as δ_1 approaches 1, the agreement $x^*=(x_1^*, x_2^*)$, approaches (1,0). In other words, as Player 1 becomes more patient, his share increases, and, in the limit, he receives all the pie. Similarly, as Player 2 becomes more patient, Player 1's share of the pie approaches zero.

Properties of the equilibrium solution

The equilibrium outcome defined above displays some important properties:

P1. (*Uniqueness*). The SPE of the game is unique, which means that the game has a determined solution.

P2. (*No delay*). Whilst the structure of the bargaining game allows negotiation to continue indefinitely, in the unique SPE agreement is reached at time t=0.

P3. (*Efficiency*). From an economic point of view, the fact that negotiation ends immediately implies that the equilibrium is efficient, in the sense that no resources are lost in delay.

P4. (*Patience*). The model predicts that when a player's discount factor increases, which means that he/she becomes more patient[6], his/her negotiated share of the pie increases. Thus, the bargaining power depends on players' relative degree of impatience.

P5. (*A-symmetry*). The structure of the alternating-offer bargaining game proposed by Rubinstein is asymmetric in one respect: one of the bargainers is the first to make an offer. This results in an advantage for the first mover who obtains, in the unique SPE, more than half of the pie. The asymmetry in the structure of the game is, however, artificial and its effects can be diminished by reducing the amount of time that elapses between periods. Rubinstein proves that, in the limiting case (i.e. when the length of the periods shrinks to 0), the amount received by a player is the same regardless of which player makes the first offer. The unique SPE then approximates the (symmetric) Nash bargaining solution.

---

[6] See footnote 3.

4 Extensions of the Standard Non-Cooperative Bargaining Model

4.1 Multiple Players

Starting from the basic Rubinstein alternating-offer game described in the previous section, most of the literature on non-cooperative bargaining theory has been devoted to models of *two* players. In many real-life situations, however, bargaining processes involve a large number of individuals or interest groups. In such a case, the prediction of the standard model that a unique equilibrium exists where agreement is reached immediately, does not usually hold. This section will discuss how the standard results change in a multilateral negotiation context, which problems may arise and what solutions have been proposed in the literature.

To simplify the discussion we consider a situation in which *three* players negotiate on the partition of a cake of size 1. There are, in fact, many ways of extending the Rubinstein two-players alternating-offer game to this case. An extension that appears to be quite natural is the one suggested and analyzed by Shaked (1986). Shaked's game is the following. In the first period, Player 1 proposes a partition $x=(x_1, x_2, x_3)$, with $x_1+x_2+x_3=1$ and Players 2 and 3 in turn accept or reject this proposal. If either of them rejects it, then play passes to the next period, in which it is Player 2's turn to propose a partition and Players 3 and 1 respond sequentially. If at least one of them rejects the proposal, then again play passes to the next period, in which Player 3 makes an offer and Players 1 and 2 respond. Players rotate in this way until a proposal is accepted by all responders. Players' preferences are represented by the utility function $u_i = \delta^{t-1}x_i$ (where $0 \leq \delta \leq 1$ is the common discount factor) and thus satisfy the assumptions A1 through A6 of the basic Rubinstein's game. Moreover, there are no exogenously imposed limits on the duration of the game, but the absence of agreement (that is bargaining forever) leads to a payoff of 0 for all players.

This model, of course, reduces to the standard alternating-offer game when there are exactly two players. Unfortunately, however, for $n \geq 3$ the game admits many equilibrium outcomes. In particular, it has been proved that: *every allocation of the cake can be sustained as a subgame perfect equilibrium (SPE) if players are sufficiently patient ($\delta > 1/2$) and outcomes with delay are also possible equilibria.* Changing the order of moves, the simultaneity of responses, etc., does not alter this conclusion.

11

The indeterminacy of the *three* (or *n*) player game has aroused much interest among researchers, and various solutions have been proposed to isolate a unique equilibrium outcome. Some authors, for example, suggested the adoption of different (more refined) equilibrium concepts, while others to modify the structure of the game.

Remaining in the context of the original unanimity model introduced by Shaked, it has been noticed that the only subgame perfect equilibrium in which players' strategies are *stationary* has a form similar to the unique SPE of the two-player game. In particular, Herrero (1985) showed that if players have time preferences with a common constant discount factor δ, then this equilibrium leads to the following division of the pie:

$$\left(\frac{1}{1+\delta+\delta^2}, \frac{\delta}{1+\delta+\delta^2}, \frac{\delta^2}{1+\delta+\delta^2} \right)$$

which tends to the equal split as δ tends to 1.

The notion of *stationary SPE* may therefore be used to restore the uniqueness of the equilibrium in multilateral bargaining situations. However, the restriction to stationary strategies is quite strong. Such strategy prescribes actions in every period that do not depend on time, or on events in previous periods. Thus, for example, a stationary strategy in which Player 1 always makes the proposal (1/2, 1/2) means that even after Player 1 has made the offer (3/4, 1/4) for a thousand times, Player 2 still believes that Player 1 will make the offer (1/2, 1/2) in the next period, which is quite unrealistic.

A more appealing way to solve the problem of indeterminacy of the *n*-player game is to modify the *structure* of the game. For example, Jun (1987) and Chae and Yang (1988, 1994) consider procedures where players are engaged in a series of *bilateral* negotiations and any player that reaches a satisfactory agreement may "exit" the game. A more interesting approach is suggested by Krishna and Serrano (1996), where players still have the possibility to exit, i.e. to leave with their share before the entire bargaining process is completed, but, unlike Jun/Chae and Yang's mechanisms, the offers are made to all players simultaneously and thus the bargaining is *multilateral*.

In particular, the structure of the game is as follows. There are three players bargaining on the partition of a 'pie' of size one. In the first period, Player 1 proposes a division $x=(x_1, x_2, x_3)$ and in response of such proposal the following situations can occur: *(a)* Both Players 2 and 3 *accept* the offer. In this case, the

game ends with that division. *(b)* Both players *reject* the offer and the game passes to the next period where Player 2 makes an offer and Players 3 and 1 must respond. *(c)* One of the responders, say Player 3, accepts the offer x while the other (Player 2) rejects. In this case, 3 can "exit" the game with an amount x_3 while players 1 and 2 are left to bargain over the division of $1 - x_3$ in period 2. The bargaining now proceeds as in the two player alternating offer game with player 2 proposing some partition of $1 - x_3$. In this model, then, the person making offer receives a payoff if and only if all the other players accept her offer, but a responder who is satisfied with her share, can simply 'take the money and run away', with no need for unanimous consensus as required in Shaked's game. With the introduction of such procedure, the authors are able to identify a unique perfect equilibrium, for any number of players. Moreover, for all *n*, the unique equilibrium is characterized exactly as in the case of two players and the equilibrium agreement approximates the *n*-player Nash bargaining solution when players are patient.

In our discussion of multilateral bargaining situations we have deliberately omitted an important element that may appear in negotiation contexts with 3 or more players, that is the possibility for players to form *coalitions*. This element makes the modelling of such situations even more difficult because one should not only determine what each player gets *individually*, but also *which coalition* will or will not form. The study of coalition formation becomes particularly important when we consider negotiations over public goods, such as many international environmental negotiations. In this case, the presence of externalities may induce players to free ride on the negotiating agreement in order to enjoy the benefits from cooperation without paying any cost. These and other problems will be discussed in Section 4 which is entirely devoted to coalition theory.

4.2 Multiple Issues

Many real-life negotiations (such as trade or environmental negotiations) do not only involve a large number of individuals, but also a set of different issues. By contrast, most of the existing literature focuses on the problem of dividing a '*single-pie'* between two agents. In this section, we will first discuss when the insights from the classical theory still apply to the multiple-issue case, and we will then consider other important elements that may emerge when players negotiate over more than one project.

13

In general, we can distinguish two different ways of handling multiple-issue negotiations. The first one is to bundle all the issues and discuss them simultaneously (*complete package approach*[7]); the second one is to negotiate the issues one by one (*sequential approach*). Suppose, for example, that there are two players, 1 and 2, negotiating, via an offer-counteroffer bargaining procedure, over two different projects, *X* and *Y*. According to the first approach, an offer is a pair (*x, y*) specifying a division on both issues, and players make offers and counteroffers of (*x, y*) until an agreement is reached. On the contrary, the second approach involves a sequential determination of allocations for the two projects. For example, players may start making offers and counteroffers on *x* only, until agreement. Once an agreement is reached, the allocation is implemented and bargaining proceeds over *y*. Intuitively, when the first approach is adopted, that is all issues are bargained simultaneously and allocations are implemented only after agreement has been reached on the whole package, then even complex negotiations reduce to "as if" single-pie bargaining and the classical theory applies directly. This conclusion is not obvious anymore when bargaining or agreement implementation take place according to the second approach, that is in a sequential way. In such a case, the order in which problems are discussed may assume a strategic role and affect the final outcome of the negotiation (in the example described above, players could start negotiating on *y* instead of *x* and thus obtain a different result).

In multiple-issue negotiations, the timing of projects on the bargaining table is specified by a *negotiation agenda,* which can be defined *exogenously*, i.e. before negotiation begins, or *endogenously*, i.e. during the bargaining process.

In general, players may have different preferences over different agendas. In the initial example, for instance, player 1 may prefer the agenda *XY* to the agenda *YX*, while player 2 may prefer *YX* to *XY*. This is because: *a)* Players may have different time constraints for reaching agreements on the two issues, that is each player may have its own *deadline* for each issue; *b)* Players may differ for their *attitude towards time*, i.e. for their discount factors. One player may, for instance, gain utility with time and have an incentive to reach a late agreement (*patient* player), while the other may lose utility with time and try to reach an early agreement (*impatient* player). In an 'issue-by-issue' bargaining process, this

---

[7] In other literatures, this is known as 'issue linkage'. The basic idea behind this mechanism is to design a negotiation framework in which players do not negotiated only on one issue,

disposition of negotiators may strongly influence strategic behaviours and, therefore, the negotiation outcome.

In the last decades, the study of the role of negotiation agenda has obtained increasing attention among researchers and various interesting contributions have been proposed in the literature. Fershtman (1990), for example, considers a situation in which two players with time preferences and additively separable utility functions negotiate, according to an alternating offer procedure, over two linear issues. In this model, the agenda is defined *exogenously* and both players are assumed to have identical discount factors and no deadlines. The author analyzes sequential agendas where the realization of utilities is postponed until both projects are accepted (simultaneous implementation). He shows that a player prefers the first project be least important to him but most important to the opponent. However, as players become increasingly patient, the impact of the agenda disappears. In and Serrano (2003) develop a model to investigate exactly the effects of agenda restrictions on the properties of the equilibrium outcome. What is found is that when the agenda is very restricted (such as, for example, when bargainers are forced to negotiate only one issue at a time, the one chosen by the proposer at each round), *multiple equilibria* and *delay* in agreement do usually arise.

In a similar setting with two linear issues and two players, Busch and Horstmann (1997) partially *'endogenize'* the bargaining agenda by introducing a separate bargaining round over it. The order in which issues are negotiated becomes, however, truly endogenous in Inderst (2000), where players bargain over projects without any ex-ante agreed agenda.

In this model issues are either mutually beneficial or strictly controversial and each subset of projects is immediately implemented after partial agreement on this set (sequential implementation). The author first derives the equilibrium payoffs when an exogenously given agenda requires that bargaining proceeds simultaneously or sequentially over the set of projects. The analysis reveals that the agenda can have a marked impact on payoffs and – in contrast to the result reported by Fershtman (1990) – this impact does not seem to vanish as players become increasingly patient. In particular, bargaining simultaneously over a set of projects

---

but force themselves to bargain on two or more issues jointly.

can improve efficiency by creating *trading opportunities* across issues[8]. Moreover, changing the agenda may have a *distributive effect*, and players may therefore prefer different agendas. In the second part of the paper the author then identifies which agenda is chosen endogenously. The results of the analysis can be summarised as follows: *A)* when issues are *mutually beneficial*, then players choose to bargain simultaneously over all issues. However, if the bargaining set contains *B) strictly controversial* projects two different sub-cases need to be distinguished depending on whether or not randomization devise is an available option: *(B1)* if players have access to a randomization device, an analogous result holds as in the previous case; *(B2)* with strictly controversial projects and without lotteries there might be multiple equilibria involving even considerable delay.

Bac and Raff (1996) focus on the effect of *incomplete information* about bargaining strength on the choice of the bargaining procedure. The model involves two players negotiating in a Rubinstein fashion over two pies, each of size one. The price-surplus is known to agents and for both players the discount factor is assumed to be equal over all issues. However, agents have asymmetric information about discounting factors. One player is perfectly informed, while the other is uncertain about his opponent's discounting factor. In particular, this can take one of the two values, δ_H with probability π, and δ_L with probability $(1-\pi)$. This bargaining game has a sequential equilibrium with rationalizing beliefs such that, while a weak (impatient) player prefers to negotiate simultaneously over the two

---

[8] The profitability and effectiveness of linkage strategies have been largely studied, especially in the literature on coalition formation. Pioneering contributions are those by Tollison and Willett (1979) and Sebenius (1983), who proposed linkage mechanisms to promote cooperation on a number of matters, such as security and international finance. Issue linkage was formally introduced into the economic literature on international environmental cooperation by Folmer et al. (1993) and by Cesar and De Zeeuw (1996) to solve the problem of asymmetries among countries. The intuition is that, if some countries gain from cooperating on a given issue whereas other countries gain from cooperating on another one, by linking the two issues it may be possible to obtain an agreement which is profitable to all countries. Linkage strategies can also be used to mitigate the problem of free-riding which normally affect negotiations over public goods, such as environmental quality. This aspect has been addressed in various ways. For instance, Barrett (1995, 1997), proposes linking environmental protection to negotiations on trade liberalisation. In this way, potential free-riders are deterred with threats of trade sanctions. Other interesting contributions are those by Carraro and Siniscalco (1995, 1997), and Katsoulacos (1997), where environmental cooperation is linked to cooperation in R&D. In a more recent work, Alesina et al. (2001) further analyse the problem of the effectiveness of linkage mechanisms in increasing cooperation, and identify an interesting trade-off between the size and the scope of a coalition: a coalition where players cooperate on too many issues may be formed by a few players, which implies small spillovers among them, whereas coalitions in which cooperation is restricted to few issues may be joined by many players, thus raising many positive externalities within the coalition.

pies, a strong (patient) player may make an offer on just one pie in order to signal bargaining strength. The uniformed player always makes a combined offer on the two pies, which may include screening the informed player and thus causing delay. According to this result, issue-by-issue negotiations may thus arise from signalling considerations.

More recently, Fatima et al. (2003) studied the strategic behaviour of agents by using an agenda that is partly exogenous and partly endogenous. This is done by decomposing the N issues into k equal stages. The issues for each stage are determined exogenously, while the order in which issues are settled at each stage is determined endogenously. The analysis shows that the negotiation outcome changes with the value of k and that the optimal number of decompositions for an agent depends on the negotiation parameters. In some negotiation scenarios the optimal value of k differs for the two agents, while in others it is identical. In other words, there exist negotiation scenarios where the utility to both agents can be improved by negotiating in stages compared to the utilities they get from single-stage negotiations.

This result complements the explanations provided by the previous works, namely that differing preferences over issues play an important role in determining negotiation agendas. Exploring the agents' strategic behaviour by separating negotiation over the agenda from negotiation over the issues can be another promising line of research.

4.3   Incomplete Information

Information can be defined as the knowledge about all factors that affect the ability of an individual to make choices in any given situation. For example, in bargaining between a *buyer* and a *seller*, information includes what an agent knows about its own parameters (like his reservation price or his preferences over possible outcomes) and what he knows about his opponent's parameters.

A critical assumption of the Rubinstein (1982)'s alternating offer game is that each player has *complete* information about the other's preferences. This assumption is quite limiting because in real bargaining there are always some parameters agents are uncertain about.

When incomplete information exists, new elements appear: a player, for instance, may try to conclude from the other player's moves who his opponents

really is; the other player, in turn, may try to bluff that he is tougher than he actually is, and so on.

An important distinction in the ambit of incomplete information models is that between *symmetric* and *asymmetric* information. Consider, for example, a game with two players. The symmetric case corresponds to the situation in which *both* players lack information about the opponent's parameters; in the asymmetric case, on the contrary, uncertainty affects just *one* of the agents.

Following Harsanyi and Selten (1972), models of games with incomplete information usually proceed by adopting the assumption that each player starts with the same probability distribution on other players' private information and that these priors are common knowledge. This is modelled by having the game begin with a probability distribution, known to all players. Thus, agents not only have priors over other players' private information, they also know what priors the other players have over their own private information.

Starting from this idea, Rubinstein (1985) proposes an extension of his original model to handle information uncertainty. This is a two persons infinite horizon game that considers incomplete information over agents' discounting factors. One of the players, say player 2, may be one of two types: weak (for high discounting factor) and strong (for low discounting factor). Player 1 adopts an initial belief about the identity of player 2. Player 1's preference is known to player 2. Agreement is reached in the first or second time period. The main result of the work is the existence of a unique sequential equilibrium when player 1's belief that player 2 is of type weak is higher than a certain threshold, and another unique equilibrium when this belief is lower than the threshold.

Within a similar framework, Fundenberg and Tirole (1983, 1985) analyse a buyer-seller infinite horizon bargaining game in which reservation prices are uncertain, but time preferences are known. In particular, they focus on whether or not the bargaining outcome can be *ex-post efficient* in the presence of one-sided and two-sided uncertainty. When exactly one player's reservation value is her private information (asymmetric case), the efficiency of the bargaining outcome depends on whether or not the players' reservation values are independent of each other. If the players' reservation values are independent, then the bargaining outcome can be ex-post efficient. If they are, instead, correlated the bargaining outcome will not be efficient. When each player's reservation value is her private

18

information (symmetric case), the bargaining outcome cannot be ex-post efficient whether or not the players' reservation values are independent of each other.

Uncertainty over agent deadlines has been studied by Sandholm and Vulkan (1999) in a symmetric information scenario. Since each player's deadline is private information, there is a disadvantage in making offers. Any offer reveals some information about the proposer's deadline, namely that it cannot be very long. If it were, the proposer would stand a good chance of being able to out-wait the opponent, and therefore would ask for a bigger portion of the surplus than it did. Similarly, the offerer knows that it offered too much if the offer gets accepted: the offerer could have done better by out-waiting the opponent. The main result of this work is that there exists a sequential equilibrium where agents do not agree to a split until the first deadline, at which time the agent with the later deadline receives the whole surplus. This result holds both for pure and mixed strategies and, in most cases, is not affected by time discounting and risk aversion.

In a more recent work, Fatima et al. (2002) address uncertainty over two parameters: deadlines and reservation prices. In contrast with the previous models, however, they assume that the probability distribution over these factors is private knowledge for each player. As in Sandholm and Vulkan (1999), the optimal strategies give the entire surplus of price to the agent with the longer deadline. However, time discounting is not neutral anymore, but affects agents' payoffs.

To conclude, it is worth mentioning a model proposed by Petrakis and Xepapadeas (1996) to study the problem of international environmental cooperation under moral hazard. This work is more related to the literature on coalition formation, but can provide interesting insights for the analysis and comprehension of bargaining processes in the presence of information uncertainty. The set of players consists of the following two groups of countries: environmentally conscious countries (ENCCs) and less environmentally conscious countries (LENCCs). The authors analyse the conditions under which the two groups can form a stable coalition to adjust emissions so that a first-best global welfare optimum is achieved[9].

The interesting aspect of the model is that asymmetries in information among countries are considered, in the sense that countries entering into agreements know their own emissions but cannot observe the emissions of the

---

[9] In particular, a self-financing side payment scheme is determined, capable of securing a stable partial coalition of ENCCs with a subset of LENCCs.

other participating countries. This may create problems in the enforcement of the agreement, since countries have an incentive to cheat by emitting more than the agreement stipulates. A mechanism that detects cheating is developed by the authors in order to induce the desired emissions even when the emission level of an individual country cannot be observed by the rest of the participating countries.

From this brief review, it seems clear that the presence of uncertainty in information may have a strong impact on the negotiation outcome and may provide an appealing explanation for bargaining inefficiencies. Informational differences may also explain the presence of bargaining power among agents and may have different effects on the negotiation outcome when players are characterized by different degrees of risk aversion.

4.4   Bargaining in stochastic environments

As discussed in the previous section, incomplete information refers to uncertainty over players' parameters, such as players' discount factors, deadlines or reservation prices. However, there are many other forms of uncertainty which may affect a bargaining process. For example, the size of the 'pie' over which agents are negotiating may vary stochastically, as well as the disagreement point. These sources of uncertainty concern the *environment* in which negotiations take place.

The theoretical literature on strategic bargaining in stochastic environment is still quite limited, as well as its applications to existing data. In the last decade, however, this issue has attracted increasing attention among researchers and various efforts have been made in this direction.

In particular, Merlo and Wilson (1995) have proposed an extension of the basic Rubinstein two-player alternating-offer game to a *K*-player bargaining model with complete information, where both the identity of the proposer and the size of the pie follow a stochastic process[10].

In each period, a state is realized which determines the cake (i.e., the set of possible utility vectors to be agreed upon in that period) and the order in which players move. The selected player may either propose an allocation or pass. If he proposes an allocation, each of the remaining players in turn accepts or rejects the

---

[10] More specifically, both random parameters follow a general Markov process, which is formally defined as a discrete process in which the probabilities of transitions from one state to another are fixed and independent of time – that is, the system at time t+1 depends only on the system at time t, and not on the state at any earlier time.

proposal. If any player rejects the proposal, a new state is realized and the process is repeated until some proposed allocation is unanimously accepted.

More formally, the model can be described as follows. Let $K= \{1,..., K\}$ denote the set of players involved in the bargaining process and let $S=\{s_0,..., s_t\}$ denote the set of possible states of the world. A *stochastic sequential bargaining game* for K may be indexed by (C, ρ, β), where for each state $s \in S$, $C(s)$ is a cake representing the set of feasible utility vectors that may be agreed upon in that state, $\rho_i(s)$ denotes the identity of the player who makes the ith move in that state, and β is the *common* discount factor for the players.

The game is played as follows. Upon the realization of a state s, $\rho_1(s)$ (the player who makes the first move in state s) chooses to either pass or propose an allocation in $C(s)$. If he proposes an allocation, player $\rho_2(s)$ responds by either accepting or rejecting the proposal and after him, all the other players respond in the order prescribed by $\rho(s)$. If the proposal is not unanimously accepted, but some players reject it, then the game moves to the next period where a new state s' is realized according to a Markov process $\sigma=(\sigma_0, \sigma_1, \sigma_2,...)$, defined in the space S. This procedure is then repeated except that the order of moves is determined by $\rho(s')$ and the proposal must lie in the set $C(s')$. The process continues until an allocation is proposed and accepted by all players.

An outcome of this bargaining game is either a pair (η, τ) – where τ denotes the period in which a proposal is accepted and η denotes the proposed allocation which is accepted in state s – or disagreement. Then, for the game starting in state s, an outcome (η, τ) implies a von Neumann-Morgenstern payoff to party i, $E[\beta^\tau \eta_i / \sigma_0 = s]$.

In order to solve the game, the authors focus on *stationary* sub-game perfect outcomes and payoffs, that is, on outcomes and payoffs generated by a stationary sub-game perfect strategy profile. The reason for this choice is that, when there are more than two players, the game does not usually admit a unique equilibrium outcome but multiple equilibria, even in the absence of uncertainty. As discussed in section 4.1, stationary is the solution concept which is typically adopted in multilateral bargaining models in order to solve the problem of indeterminacy of the negotiation outcome.

The main results of the paper can be summarized as follows:

R1: There exist a *unique* (stationary sub-game perfect) equilibrium;

R2: The equilibrium is *efficient*, even though it may involve *delays*.

This result is not exactly conforming to what the standard literature predicts. In particular, according to the traditional models of bargaining, when an equilibrium exists, either it is efficient and such that agreement is reached immediately (as in the basic Rubinstein two-player game), or outcomes with delay may arise but efficiency is not guaranteed anymore.

In the standard theory, the most common explanation for delaying agreement is that players are unsure about the true preferences of their opponents. In other words, *incompleteness* in information (see section 4.3) can cause inefficient equilibrium outcomes. In the context of complete information, sequential bargaining models generally admit delays only if there are *multiple equilibria*.

On the contrary, in the stochastic model by Merlo and Wilson, which is a model with *complete information*, agreement may be delayed even in the *unique* stationary sub-game perfect equilibrium and the equilibrium is still efficient.

The intuition for this result is that, when the future size of the cake is random, there can be potential benefits to waiting as the size of the cake may grow in the future. In other words, delay is caused by the expectation that the total bargaining value may rise in the future and hence is efficient from the point of view of the negotiating parties.

Various applications exist of the framework described above, which mainly focus on the problem of government formation. Merlo (1997), for instance, investigates the process of government formation in post-war Italy, while Diermeier et al. (2004) explore the role of bicameralism in determining government durability.

These studies seem to confirm the efficiency of delays predicted by Merlo and Wilson for bargaining in stochastic environments. From a theoretical point of view, however, this result depends also on other features of the game, such as the agreement rule which is adopted, or the bargaining procedure. For example, when the agreement rule is a general *q*-quota rule as in Eraslan and Merlo (2002) , uniqueness and efficiency of the equilibrium are not guaranteed anymore. On the other hand, when players are given the possibility to delay making offer, as in Furasawa and Wen (2001), the game still has a unique equilibrium solution, but outcomes with delay are not efficient.

To conclude, it is important to notice that all the models mentioned above are static models of bargaining, while the problem of negotiating over a pie of not-

fixed size could be better seen as a dynamic problem, where the state of the system evolves in time. In non-cooperative bargaining theory, the dynamic nature of negotiations is normally represented trough repeated bargaining games, which are described in the following section. Further research is therefore needed in order to understand what are the links between stochastic and dynamic nature of the bargaining setting, and how this can be modelled.

4.5   Repeated Bargaining Situations

An implicit assumption of the Rubinstein (1982)'s bargaining model is that players' interaction ceases after a decision is reached, in other words, once the negotiation process ends, players do not meet anymore. In fact, this is rarely the case in real settings because agents usually have the opportunity to be involved in a *sequence* of bargaining situations. Think, for example, of two adjacent countries and of the vast occasions of bargaining they may have over time: from trade to international protection, from political questions to environmental problems…

This section will focus on *repeated bargaining games*, which have been proposed in the literature in the attempt to represent the long-term relationships that may exist among bargainers. In a repeated framework, a game is played in successive stages and at each stage players can decide on the basis of the actions and the outcomes of the previous stages. There is an accumulation of information about the 'history' of the game that may affect players' strategic choices. In particular, even if it is the 'same game' which is repeated over a number of periods, the global 'repeated game' becomes a fully dynamic system with a much more complex structure than the one-stage game.

We will analyse here a simple repeated bargaining situation in which two players sequentially bargain over the partition of an infinite number of cakes. The model is based upon Muthoo (1995) and consists in an infinite repetition of the standard Rubinstein model. Despite its simplicity, it provides some interesting results and allows us to lay down the basic structure of repeated bargaining situations. The first important qualification of the game is that players start bargaining over the partition of the (n+1)th cake (where n=1, 2,…) if and only if they reach agreement on the partition of the nth. The second qualification is that the time at which the players start bargaining over the partition of the (n+1)th cake is determined by the time at which agreement is struck over the partition of the nth cake. The structure of the game is as follows: there are two agents A and B, who

23

bargain over the partition of a cake of size π ($\pi > 0$), according to an alternating-offer procedure. If agreement is reached at time t_1, then immediately the players consume their respective (agreed) shares. Then τ ($\tau > 0$) time units later, at time $t_2 = t_1 + \tau$, the players bargain over the partition of a second cake of size π. Agreement at time t_2 is followed immediately with players consuming their agreed shares. This process continues indefinitely (t_3, t_4,…), provided that players always reach agreement. However, if players perpetually disagree over the partition of some cakes, then there is no further bargaining over new cakes: agents have simply terminated their relationship.

In this model, the payoffs to the players depend on the number N of cakes that they partition. In particular, if N=0 – that is they perpetually disagree over the division of the first cake – then each player's payoff is zero. If N >0, then player i's payoff is:

$$\sum_{n=1}^{N} x_i^n \delta^{t_n} = \sum_{n=1}^{N} x_i^n \exp(-r_i t_n)$$

where x_i^n is player i's share of the nth cake, t_n is the time at which agreement over the partition of the nth cake is struck, and δ_i is player i's discount factor, with $\delta_i^{t_n} = \exp(-r_i t_n)$.

This game has a unique stationary subgame perfect equilibrium and in equilibrium agreement is reached immediately over the partition of each and every cake. In general, however, this equilibrium outcome is different from the unique SPE partition of the *single* available cake in Rubinstein's model. The intuition for this difference is as follows: in a repeated bargaining model a player's discount factor determines not only her cost of rejecting an offer, but also her value of future bargaining situations. Suppose that player i becomes more patient (that is her discount rate r_i decreases and her discount factor δ_i increases). This means that her cost of rejecting an offer decreases. However, it also means that her value of future bargaining situations increases. When bargaining over the partition of a cake, the former effect increases her bargaining power (as she is more willing to reject offers), but the latter effect decreases her bargaining power because she is more willing to accept offers so that the players can proceed to bargain over the partition of the next cake. It has been shown that, under some plausible conditions, the latter effect tends to dominate the former effect. This result implies that when a player becomes less patient, she receives a greater share of each and every cake. Thus, the

impact of players' discount rates in repeated bargaining situations may differ fundamentally from that in one-shot bargaining situations.

4.6 Synthesis of the results

Table 1 in the following page summarizes the results of the analysis conducted in sections 2 and 3. In particular, three main characteristics of the equilibrium outcome are considered for the basic Rubinstein alternating-offer game and its extensions: *(i)* the determinacy of the equilibrium, *(ii)* the timing of the agreement and *(iii)* the efficiency of the result.

As previously noted, the model proposed by Rubinstein involves only *two* players bargaining over the division of a *single* 'pie' in a *complete information* setting. Under these conditions, the alternating-offer bargaining game admits a *unique SPE*. In such equilibrium, the agreement is reached *immediately* and the bargaining process is *efficient*, in the sense that no resources are lost in delay.

Table 1: Characteristics of the equilibrium outcome in the basic Rubinstein model and in the extensions analysed through section 3.

	Determinacy of the equilibrium	Timing of the agreement	Efficiency of the equilibrium
Rubinstein (1982)	Unique SPE	No delay	Guaranteed
Multiple players	Multiple equilibria	Possibility of delay	Non-guaranteed.
Multiple issues *A) Mutually beneficial issues*	Unique equilibrium	No delay	Guaranteed
B) Strictly controversial issues: *B1) possibility of randomising*	Unique equilibrium	No delay	Guaranteed
B2) no lotteries	Multiple equilibria	Possibility of delay	Non-guaranteed
Incomplete information	Unique equilibrium	Presence of delay	Not always guaranteed

Bargaining in a stochastic environment (*uncertainty about the size of the pie and/or the order of moves*) • *agreement rule unanimity*	Unique equilibrium	Possibility of delay	The equilibrium is efficient even with delay
• *agreement rule different than unanimity*	Unique equilibrium	Possibility of delay	Not normally guaranteed
Repeated bargaining games	Unique equilibrium*	No delay	Not always guaranteed

*In general, however, this equilibrium outcome is different from the unique SPE of the basic Rubinstein's model because of the different impact that players' discount rates usually have in repeated bargaining situations. In particular, in the standard model, as a player becomes more patient, her share of the pie increases, while the opposite happens in repeated bargaining games (see section 3.4.).

To which features of the model can we attribute this result? As shown in Table 1, one possible explanation for delaying agreement is that players are unsure about the true preferences of their opponents. In other words, *incompleteness* in information can cause inefficient equilibrium outcomes.

In a complete information setting, the presence of delay is closely related to the existence of *multiple equilibria*, which may arise, for instance, when the negotiation process involves more than two players. However, if the negotiation takes place in a stochastic environment (such as, for example, when the size of the pie over which players bargain varies stochastically) agreement may be delayed even in the *unique* sub-game perfect equilibrium and the equilibrium may still be efficient[11].

---

[11] Another important element that may affect the timing of the solution is the presence of option values, which do normally arise in dynamic contexts, that is when the state of the system evolves in time. This aspect has been widely studied in optimal control theory, while research is still needed in the ambit of non-cooperative game theory.

5   Non-Cooperative Coalition Theory

As noticed in Section 4.1, many real negotiations involve a large number of parties or interest groups. When modelling these situations, the standard bargaining theory (both cooperative and non-cooperative) makes the implicit assumption that there are just two possible outcomes of the bargaining process: the *cooperative outcome*, where an agreement is reached among *all* players involved in the negotiation process, and the *non-cooperative outcome*, where no agreement forms. This dichotomy is often not representative of real-life situations where *partial agreements* can form among a subset of players.

In this section, we will focus on a different approach to *multilateral* negotiations, which is 'Non-cooperative Coalition Theory' (NCT). Unlike standard bargaining theory, this approach is able to take into account these *intermediate* cases because it allows for the possibility of sub-coalitions to form.

More in general, we can distinguish both a cooperative and a non-cooperative perspective within the theory of coalitions. However, as we will see in section 4.1, cooperative coalition theory (CCT) basically coincides with the standard cooperative bargaining theory (Nash, 1950) for the case of *N* players and it cannot really help in understanding the forces which drive the formation of (partial) coalitions.

As shown in some recent works by Gomes and Bloch, another reason to concentrate on NCT is that this is, in general, more suitable to analyse the problem of coalition formation within the context of negotiations, because it is more focused on players' incentives to cooperate and on the procedures which lead to the formation of coalitions.

5.1   Cooperative versus non-cooperative coalition theory

As emphasized by Bloch (1997), the analysis of endogenous formation of coalitions poses three basic questions: (1) Which coalitions will form in equilibrium? (2) How will the coalitional worth be divided among coalition members? (3) How does the presence of other coalitions affect the incentives to cooperate?

The *cooperative approach* to coalition formation mostly focuses on the second question, that is the division of the payoffs among co-operators, while the first question is generally avoided. Most of cooperative models, indeed, are based

27

on the idea that, among all the possible coalitions that could form, the one that is most valuable[12] will actually be produced. Therefore, there is an assumption of Pareto-optimality (i.e., the most efficient, value-maximizing coalition will always form) regardless of the process required to form such a coalition. In fact, the processes are considered unproblematic, as rational actors will always choose the outcome that maximizes their value. In other words, once the best outcome is determined based on the attributes of the actors and the payoffs available to them, the assumption is that the outcome will *always* realize.

This is exactly the same idea of Nash/cooperative bargaining theory which imposes a number of axioms on the bargaining solution and, assuming that all players participate in the agreement, focuses on the problem of dividing the pie according to some criteria (such as feasibility, fairness, stability). The solution concepts adopted are the same: from the Core to the Shapley-value, the Nucleolus, the Kalai-Smorodinsky solution…

The third question, dealing with competition between coalitions, is simply ignored in traditional cooperative coalition theory (as well as in cooperative bargaining theory, where competition among players is not really taken into account). The analysis is based, indeed, on the *characteristic function* that assigns to each coalition C a real number $v(c)$ representing the worth of the coalition. The worth, however, is defined as the aggregate payoff that a coalition can secure for itself *irrespective* of the behaviour of players outside the coalition. Then spillovers between coalitions are not allowed.

Because of these limitations, cooperative games, which were prevalent in earlier coalition theory literature, have largely given way to non-cooperative games of coalition formation. The *non-cooperative* approach is based on the *partition function* that assigns an individual payoff to each player for each possible coalition structure. This is a generalisation of characteristic function games that allows for considerations of spillovers. In particular, if the worth of a coalition C is independent of the coalitions formed by the other players, the two definitions coincide. If, on the other hand, the formation of coalitions affects all the players in the game, there is no univocal relationship between partition functions and characteristic functions, and a game in partition function form carries more information than a game in characteristic function form.

---

[12] Value is not usually defined explicitly, but is assumed to have some material weight. An example might be the amount of policy power the winning coalition possesses.

In general, a non-cooperative game of coalition formation can be modelled as a two-stage game: in the *first stage*, players decide non cooperatively whether or not to join a coalition given the adopted burden-sharing rule; in the *second stage*, agents set their policy/decision variables by maximizing their welfare function given the decision taken in the first stage and the adopted burden-sharing rule. The standard assumption is that coalition members act as a single player maximising the aggregate payoff to their coalition, but behave non-cooperatively towards outsiders. Equilibrium coalition structures are then determined by applying the concept of internal and external stability (Barrett 1994, 1997; Carraro and Siniscalco 1993; Hoel 1992; Hoel and Schneider 1997; Rubio and Ulph, 2001). Internal stability means that no coalition member has an incentive to leave its coalition to become a singleton, and external stability that no singleton has an incentive to join a coalition, assuming that the remaining players do not revise their membership decision[13].

With this simple framework non-cooperative coalition theory can capture players' incentive to cooperate without the need to make assumptions on the set of possible outcomes, as standard bargaining theory does.

The two-stage approach described above represents the common denominator of non-cooperative models of coalition formation. However, such models may differ substantially with respect to other important features: the order of moves, the membership rules, the players' conjectures, the type of free-riding in games with spillovers, and so on. By changing these features of the game, the final coalition structure changes.

5.2   Simultaneous (non-cooperative) games

A first important distinction is that between *simultaneous* and *sequential* (non-cooperative) games of coalition formation. In simultaneous games, all players announce at the same time their decision to form coalitions. In such games, it appears that the set of Nash equilibria is often quite large, forcing researchers to use some refinements in order to make interesting predictions. As noticed by Bloch (1997), these refinements are usually of a cooperative nature; hence, the study of

---

[13] Most of the existing contributions restrict coalition formation to a single coalition, allowing to group players into signatories and non-signatories. However, there have been some recent developments that admit the co-existence of multiple coalitions. These approaches invoke stability concepts that consider not only deviations by single players but also by subgroups of players.

simultaneous games of coalition formation is at the frontiers between cooperative and non-cooperative game theory.

The problem of simultaneous formation of coalitions has been analysed in the literature under different coalition formation rules. Looking at the existing models, the following three membership rules can be identified: *(i) Open Membership*, *(ii) Exclusive Membership*, and *(iii) Coalition Unanimity rules*. A key difference between them lies in what can happen to the membership of a coalition once it is formed: Can an existing coalition break apart, admit new members or merge with other coalitions?

Open membership is the rule originally adopted in the literature on cartel formation (D'Aspremont et al., 1983) and in the environmental literature on international agreements (Hoel, 1992; Carraro and Siniscalco, 1993; Barret 1994). In *open membership games*, any player is free to join or leave a coalition. Accordingly, players cannot specify in advance the coalition they wish to form. Rather, they announce a message (for example, their willingness to participate in a coalition), and coalitions are formed by all players who make the same announcement.

In *exclusive membership games* (Yi and Shin, 1994) or *game Δ* (Hart and Kurz, 1983), each player can join a coalition only with the consensus of the existing members, but she is free to leave the coalition. In this decision process, each player's message consists in a list of players with whom she wants to form a coalition. Those who announce the same list will then form a coalition, which is not, however, necessarily formed by all players in the list.

Finally, in *coalition unanimity games* (Yi and Shin, 1994; Chander and Tulkens, 1997; Bloch, 1997) or *game Γ* (Hart and Kurz, 1983), no coalition can form without the unanimous consensus of its members. This implies that players are not free to either join the coalition or to leave it. Therefore, this membership rule introduces restrictions both on entry (as the exclusive membership rule) and on exit behaviours of players. In the decision process, players' messages consist in a list of players as in the previous one. However, if a coalition is formed, it is necessarily composed of all players in the list and as soon as a player defects the coalition breaks up into singletons.

Yi (1997) provides an interesting analysis of the results of simultaneous games of coalition formation for the different membership rules described above. In particular, the author considers games where the formation of coalitions creates

30

externalities on non-members (which is often the case in real economic contexts) and recognizes in the *sign* of the externalities a determinant organizing principle. In general, coalition formation may create either *positive* or *negative* externalities on outside members/coalitions. Examples of positive externalities include output cartels in oligopoly and coalitions formed to provide public goods (such as environmental quality). Examples of negative externalities are research coalitions with complementary research assets and customs unions in international trade. The main results of the analysis can be summarized as follows:

> *R1.* With *negative externalities*, and under some reasonable conditions on the partition function, the *grand coalition* is an equilibrium outcome under the Open membership rule, but typically not under Exclusive Membership and Coalition Unanimity.
>
> *R2.* With *positive externalities*, the grand coalition is rarely an equilibrium outcome for any of the membership rules mentioned above and only *partial agreements* form. The grand coalition is more likely to emerge at the equilibrium under Coalition Unanimity.

The explanation for these results is quite intuitive. If externalities are negative, there is a disadvantage for players to stay outside the coalition and then it is more likely that full cooperation is reached. On the contrary, if externalities are positive, players who do not enter into the coalition may still enjoy (part of or all) the benefits from cooperation without paying any cost. This produces incentives to free ride that, in turn, prevent the formation of the grand coalition.

> *R3.* In the presence of positive externalities, not only the grand coalition rarely forms, but also the *size* of the partial agreement(s) which arise in equilibrium is usually very *small.*

An important implication of these results is that standard bargaining theory may not be appropriate in the presence of positive externalities where the emerging equilibria are usually very far from full cooperation. Therefore, we can re-state that standard bargaining theory can be appropriately applied to negotiations among $n>2$ players only in the absence of externalities. With negative externalities, standard bargaining theory is not appropriate, but results are equivalent to those obtained by

using non-cooperative game theory. With positive externalities, the only appropriate tool is non-cooperative coalition theory.

The study of simultaneous games of coalition formation has, however, revealed a number of difficulties which is important to underline. First of all, these games do not usually admit a unique equilibrium outcome. The *multiplicity of the equilibria* imposes the use of more refined solution concepts in order to obtain a sharp prediction about the final coalition structure. Yi and Shin (1994) and Hart and Kurz (1983), for instance, propose to consider cooperative refinements such as coalition-proof Nash equilibrium and strong Nash equilibrium. These selection mechanisms, however, are in general very stringent and this might generate unrealistic predictions on the final coalition structure. Another limit of the simultaneous approach is that it does not allow to *identify* the members of a coalition because all players have to decide at the same time whether or not to participate. The identity of the players may instead be relevant for the determination of the final equilibrium outcome. Finally, in simultaneous games, players cannot be *'farsighted'* in the sense that individual deviations cannot be countered by subsequent moves. Consider, for example, the departure of a player from a coalition. In a simultaneous game, either the other coalition members remain together (in open membership and in the game Δ) or the coalition breaks apart (in the game Γ). But in both these formulations, members of the coalition which are left by the deviator are not allowed to react to the move of the deviator.

5.3   Sequential (non-cooperative) games

The problems of simultaneous games have led to the formulation of *sequential games of coalition formation* where the process is described by an explicit extensive form non-cooperative game. In the context of games *without* spillovers, sequential processes have been proposed by Selten (1981), Chatterjee et al. (1993), Moldovanu (1992) and Perry and Reny (1994), among others. In most of these games, the basic structure is an extension to *n* players of the Rubinstein's (1982) alternating-offer bargaining model described in section 2. This structure was extended to games *with* spillovers by Bloch (1996) and Ray and Vohra (1996).

All these works, although different with respect to the presence of externalities, are based on a common assumption, which is: once a coalition has been formed, the game is only played among the remaining players. The typical

structure of the game is as follows. Players are ordered according to a fixed rule[14] and the first player starts by proposing the formation of a coalition *C* to which she belongs. Each prospective member responds to the proposal in the order determined by the fixed rule. *If* one of the players rejects the proposal, she must make a counteroffer and propose a coalition *C'* to which she belongs. *If*, instead, all proposed members accept, the coalition *C* is formed. All players belonging to *C* then withdraw from the game, and the first player in *N\C* starts making a proposal.

However, the assumption of *immediate exit* usually results in inefficient outcomes, as shown in the following example inspired by Chatterjee et al. (1993). Let *n*=3 and the gains from cooperation be represented by a coalitional function *v(C)=0* if *C=1*, *v(C)=3* if *C=2*, and *v(C)=4* when *C=3*. As players' discount factor, *δ*, converges to 1, the outcome of the bargaining procedure where the grand coalition forms should result in equal sharing of the coalitional surplus among the symmetric players (4/3 for every player). But clearly, players then have an incentive to deviate forming an inefficient coalition of size 2, which induces a payoff of 3/2 for each coalitional member. If this coalition must leave the negotiation after its formation, the additional surplus of 1 is lost.

In order to avoid these inefficiencies, other authors have proposed coalitional bargaining models where agents cannot choose to exit, but they are given the possibility to renegotiated over the formation of a coalition. In particular, Seidmann and Winter (1998) have focused on games without externalities, while Gomes (2001) has extended the analysis to the case of positive and negative spillovers. In these games with *continuous renegotiations*, the grand coalition is ultimately formed, as players carry on bargaining until all gains from cooperation are exhausted. However, delays may arise in the enrichment of the agreement.

Unlike games with immediate exit, the models with continuous renegotiations do usually produce efficient equilibrium outcomes.

5.4   Coalition Formation *and* Negotiations

Both the approaches described in the previous sections do not explicitly address an important question, that is, when the members of a coalition would *voluntarily* choose to leave the negotiation table. Many real-life situations seem to suggest that

---

[14] Okada (1996) proposes a model without externalities where players are randomly selected instead of being ordered according to a fixed rule. Montero (1999) adopts a similar structure but allowing for the existence of spillovers.

this decision is a strategic action as much as the choice of forming coalitions. The Kyoto protocol to reduce the emissions of greenhouse gases, for instance, shows that countries often adopt this kind of strategies in the attempt to modify the final outcome of the negotiation.

For the first time, these problems have been addressed in the literature in a work by Bloch and Gomes (2003) where players are engaged in two parallel interactions: they propose to form coalitions in order to extract gains from cooperation; and coalitions participate in a repeated normal form game, where they choose *endogenously* when to leave the negotiation process.

More precisely, the game, which is an infinite horizon *N*-player game, is characterized by two distinct phases at every period. In the first phase, or *contracting phase*, a player is chosen randomly to propose a coalition and a payment to all other coalition members. Prospective members respond in turn to the offer and the coalition is formed only if all its members agree to the contract. If a coalition is formed, the proposer acquires control rights over the resources of coalition members (the proposer player is then identified with the formed coalition). In the second phase, or *action phase*, all proposer players choose an action, which may be a permanent action (in which case the coalition they 'control' exits the game) or a temporary action. The action profile determines a flow payoff for all players, representing the underlying economic opportunities. The interplay between the contracting and action phases enables the authors to consider simultaneously issues of coalition formation, externalities and endogenous exit decisions.

A key feature of the model is the existence of (pure) outside options for players involved in the negotiation process. In classical two-player bargaining games, when an agent chooses her outside option, negotiations end and the other player is left with a fixed payoff. In multilateral negotiation contexts, when a player opts out and chooses to enforce a permanent action, the other players continue to bargain over the formation of coalitions and continue to choose actions which may affect the payoff of the exiting player. The authors point out that there is a crucial distinction between situations where outside option values are independent of the action of other players *(pure outside options)* and situations where players' outside option values are affected by the actions of remaining players.

The main result of the paper is that *there always exist an efficient equilibrium outcome in games with pure outside options.*. The intuition for this result is as follows. Early exit normally results in an aggregate efficiency loss. In a game with pure outside options, players are able to capture this inefficiency loss and will never choose to leave before the grand coalition is formed. By staying in the game one more period, indeed, a player is guaranteed to obtain her outside option (which remain available because outside options are pure), and is able to capture the inefficiency loss by proposing to form the grand coalition when she is recognized to make an offer. Hence, early exit will never occur in equilibrium.

The authors also provide some examples of games where the outside options are *not* pure. They show that, in such cases, the equilibrium outcomes may lead to the inefficient formation of partial coalitions. This result highlights the difference and the importance of this model with respect to the coalitional bargaining models previously mentioned. In a setting with externalities, for instance, Ray and Vohra (1999) showed that when players cannot renegotiate, the outcome of coalition formation is typically inefficient, as players have an incentive to leave the game before extracting the entire surplus. On the contrary, Gomes (2001) established that when renegotiation occurs and players cannot choose to exit, the outcome is always efficient. Bloch and Gomes (2003) identify a new type of friction – externalities on players' endogenous outside options – that may lead to bargaining inefficiencies.

6   Fair-division theory

Starting from the basic Rubinstein's alternating-offer game, almost all economic models of bargaining have remained faithful to the traditional assumptions about agent behaviour underlying the economic science, that are: *perfect rationality* and *purely selfish pursuit of personal interests*. In other words, standard bargaining theory assumes that when deciding whether or not to accept an offer, each bargainer focuses exclusively on her own payoff and compares what she can get by accepting the proposal with what she could get by rejecting it and moving to next period. According to this framework, agents do not have any fairness concern, in the sense that they do not care about the *distribution* of payoffs or the *intentions* of the other bargainers.

Yet everyday experience indicates that fairness consideration may have a significant influence on people's behaviour, and that humans are inclined to retaliate against those who treat them unfairly.

6.1 Experimental Evidence

Traditional assumptions of perfectly rational and self-oriented agents do normally work very well in the context of 'almost' perfectly competitive markets, where the number of players is 'very big' and what really matters for the economic science is the representative agent, i.e. an imaginary agent whose every single trait of character is the average of that trait over all agents present in the market (see, for instance, the experimental work by Roth et al, 1991). On the contrary, several experimental studies of bargaining situations have revealed the importance of fairness considerations in negotiation contexts. Negotiations are indeed a very peculiar type of economic interactions, because in their case the assumptions of perfect competition and 'large numbers' are inappropriate.

Most of the existing bargaining experiments examine one-period (or "ultimatum") games. In such games, the Proposer makes a take-it-or-leave-it offer to the respondent on how they should split a surplus of a fixed size. The bargaining proceeds as follows: the proposer offers a share s to the respondent (with a share $1-s$ going to herself); the offer can be accepted – in which case the respondent gets a payoff of $x_2=s$, and the proposer gets a payoff of $x_1=1-s$; or it can be rejected, in which case both players get a payoff of 0. The standard model predicts that the unique sub-game perfect equilibrium for this game is for the proposer to offer $s=0$, which is accepted by the respondent. This outcome is Pareto efficient, but it is clearly highly unequal.

The data generated by ultimatum experiments with complete information indicate that rather than making offers where the proposers keep the entire surplus minus the smallest unit of account, the proposers offer distributions that are closer to an equal split of the surplus (see, for instance, the experimental results of Thaler (1988); Güth and Tietz (1990); Roth (1995), Slonim and Roth (1997), and Ochs and Roth (1989); Spegel et al. (1990)).

Regularities and robust facts emerging from experimental studies of this type are: (i) there are virtually no offers above 0.5; (ii) most of the offers falls within the 0.4-0.5 interval; (iii) there are almost no offers below 0.2; (iv) low offers are usually rejected, with the probability of rejection being inversely related to s.

Similar experimental results are obtained for "dictator" games and public good distribution games. In the former, a dictator has to decide what share s of a given surplus he should give to his opponent: whereas the standard model predicts that $s=0$, experimental evidence indicate that around half of the subjects choose $0<s<0.5$ (see, for instance, Forsythe et al., 1988, and Andreoni and Miller, 1996). In public good contribution games, where the theoretical models predict that, because of the presence of externalities and free-riding opportunities, there will be an under (over) provision of the good (bad), experimental evidence indicates that players act cooperatively, if the possibility of punishing free-riders is introduced.

These studies suggest that it is very important to incorporate fairness in bargaining theory because this can markedly change the predictions of the models – alternatively, failure to do so means that theoretical models are of little help in predicting what the outcome of a negotiation process will be. Several attempts have been made to explain the observed persistent deviations from the theoretical predictions, based on different assumptions over the motivational structure of players.

6.2   Theories of fair behaviour

Recent studies have yielded two competing theories to explain these stylised facts. Gueth and Huck (1997), Kravitz and Gunto (1992) and Rabin (1993) distinguish between these two theories by noticing that the decision to make fair offers can be the result of two possible scenarios: (1) self-interested proposers make fair offers because they fear that unequal offers might be rejected, and (2) proposers make fair offers simply because they are motivated by fairness concerns. In the latter, *normative*, hypothesis, fair outcomes are the result of purely altruistic behaviour. In the theory developed by Rabin (1993), for instance, fairness rests on the idea of reciprocity – people want to help those who help them, and hurt those who hurt them[15] – and the notion of equity is based on the perceived *intentions* of the opponent. In the *expected utility hypothesis*, "fair behaviour" emerges out of self-interest, rather than out of altruism. Recent experiments (see, for instance, Harrison and McCabe, and Straub and Murnigham, 1995) seem to indicate that outcomes are more equitable than the theoretical models would predict not out of aversion to inequality, but rather offerers want to *appear* to be fair out of self-

---

[15] Strategy which, in repeated and evolutionary games, is called tit-for-tat – or, in a stronger sense, trigger strategy.

interest. Pillutla and Murnigham (2003) report that this can be attributed to the information structure of the game: in ultimatum games when receivers did have information on the surplus, offerers offered significantly more – they were therefore being *strategically* fair, rather than *truly* fair.

Empirical evidence over inequality preferences is conflicting: there are situations in which the standard self-interest model is rejected (games of dictator, ultimatum, and public good contribution with punishment); yet, there are also situations in which experimental evidence supports theoretical predictions (market games and public good games without punishment). There are some attempts at developing theoretical models consistent with the observed "fair" behaviour: for instance, Fehr *et al* (1999) show that this contradicting evidence can be reconciled in a unified theoretical framework if, in addition to selfish players, there is a *fraction* of players who cares about inequality. The question of what produces "fairer" than expected outcomes and behaviours remains, however, open.

6.3   Fair Division Procedures

Whatever the reason behind the emergence of "fair" outcomes and behaviours, it is now accepted that, at least under some circumstances, and at least for a fraction of the players, equity matters. The concept of fairness is subjective, and so is the reference point to which outcomes are compared, which is the result of many variables (social context, background, entitlements…). In addition, there is substantial experimental evidence indicating that nearly all subjects show aversion towards disadvantageous inequality, but aversion to advantageous inequality is much less prominent (Loewenstein et al., 1989).

These aspects are important when modelling the negotiation process: the *perception* of fairness plays a crucial role in determining how a surplus is divided, and the potential allocation rules must be perceived as "equitable" and "envy-free"[16] by all parties. This is especially true in a non-cooperative setting, where binding agreements cannot be imposed: whereas cooperative game theory can postulate abstract properties that an allocation scheme should have (axioms), negotiation theory needs to specify how, *constructively*, an allocation with the desired properties can be produced, and which strategies players should follow to

---

[16] Equitability refers to an external comparison of utilities (is my announced valuation of the goods I have received equal to my opponents?). Envy-freeness, on the other hand, is

ensure this outcome. The former approach focuses on distributive fairness (that is, the properties of the allocation scheme between interdependent individuals), whereas the latter addresses also the issue of procedural or motivational fairness, i.e. the procedures used to arrive at that allocation scheme.[17]

According to Brams and Taylor (1995), bargaining theories have proved inapplicable to the settlements of real life disputes also because of their divorce from theories of *fair division*. An allocation procedure is *fair* to the degree that it satisfies certain desirable properties, and it enables each player to achieve a certain level of satisfaction. Desirable properties for a fair allocation procedure are: proportionality; envy-freeness; efficiency; equitability; and invulnerability to manipulation.

Much of the efforts have focused on the *efficiency* of allocation procedures, with little attention to the concept of *equity*, and even less to the issue of minimising *envy*: yet, in the context of non-cooperative bargaining, self-enforcement requires that the resulting allocation of the surplus be equitable and envy-free, if an agreement is to be achieved at all. In fact, envy-free divisions are rarely Pareto-efficient[18], but, if the price of obtaining a bigger share of the surplus is that you envy somebody else, then the efficient allocation may not be feasible.

There are numerous fair division procedures, which can be classified according to the number of players they are applicable to; the properties they satisfy; the type of good/issue they are applicable to[19]. Many procedures are quite involved, and may therefore be of little use in practice – but they all rest on the idea of eliminating envy by creating ties, which is a very powerful idea. We will therefore concentrate on those refinements which may lead to proportional, envy-free and efficient allocation of a resource.

6.3.1 Basic allocation procedures

---

based on an internal comparison (would I be better off with my opponents' allocation than I am with mine?).

[17] See Beersma and Dreu, 2003, p.220.

[18] Recall that Pareto efficiency requires an allocation to be such that no improvement to at least one player could be made, without making another player worse off. A (0,1) split may therefore be Pareto efficient, but it is not likely to be envy free. Whereas a (0.5,0.5) split-which *could* be envy-free – may not be Pareto efficient, depending on the initial allocation of goods and/or the distribution of preferences.

[19] For a detailed overview, see Brams and Taylor, 1996. A branch of the literature concentrates on proving the existence of envy-free allocations, but provides no indication on how the allocation can be achieved (see, for instance, Brams and Taylor (1995) and, more recently, Marteens et al. (2002) and reference therein).

Fair allocation procedures are of two main types: (A) continuous, *moving knife* procedures, in which a mediator (or referee) proposes to agents continuous partitions of the surplus, *s*, which are strictly increasing; at any point, a player can stop the referee, and get *s* – with 1-*s* going to the other player. If the players do not know each other preferences, then they will stop the referee when, in their evaluation, *s* = 0.5. Variants of the basic moving knife procedure include the one proposed by Levmore-Cook for 3 persons; Brams and Taylor generalised trimming procedure; and Webb moving knife procedure. (B) Discrete, divide and choose type, in which one player (the divider) cuts the cake into pieces s/he values the same, and the other player (the chooser) selects one of the pieces. These procedures assure each player a piece of cake perceived to be at least 1/n of the total surplus (proportionality), no matter what the other players do; the cutter must play "conservatively" by dividing the surplus into exactly 1/n (according to his or her evaluation). Divide and choose procedures can be applied to a divisible, heterogeneous good and, for two players, they lead to envy-free allocations. Variants include Fink lone-chooser procedure; filter-and-choose; the Steinhaus-Kuhn[20] lone-divider procedure; the Banach-Knaster last diminisher procedure; and the Selfridge-Conway discrete procedure.

Both the approaches share the same characteristics of producing *proportional* allocations. Unfortunately, proportional allocation algorithms are, generally, not efficient in the economic sense. One reason for this inefficiency is that players, when choosing strategies that ensure them at least proportionality, have to forgo strategies that would give them more – for instance, divide-and-choose requires that the dividers equalise portions, even though they may prefer different parts than the choosers. To lessen this problem, pre-play communication to discover opponents' preferences could be helpful. In fact, and as shall become clearer during the discussion, the procedures described can be applied when agents are heterogeneous: it is in fact the information structure which will determine the properties of the allocation, more than the preference structure.

In the case of two homogenous players, proportionality is equivalent to envy-freeness, but this result does not extend to the case of more than two players. A procedure which is envy-free for *n*>2 players needs to be such that there exists a strategy for each player, which guarantees him or her a piece of the surplus that

---

[20] Steenhaus first proposed the method for n=3 player – which also applies to n=4 players. Kuhn (1967) provides a generalisation of the algorithm for any number of players.

s/he considers equal to the largest, no matter what the other players do. None of the *n*-person proportional procedures is envy-free: whilst they guarantee each player a portion which is at least $1/n$, one or more of the players may think that another player received a larger piece.

In addition to not being envy-free in the case of *n*>2 players, proportional fair allocation procedures cannot easily be employed in the case of non-divisible goods. In this case, the main allocation procedures is the Knaster's sealed bids procedure, an *n*-person auction scheme that is proportional and efficient, but not envy-free for *n*>2. Players submit sealed bids for items, which are then allocated to the highest offerer; in the second stage, there are some side-payments, with the monetary reallocation estimated by computing the "fair share" for each player.

6.3.2 *Refinements of the basic procedures*

Generally, there is no fair-division scheme that is simultaneously (1) algorithmic; (2) proportional; and (3) efficient. An exception is the Adjusted Winner procedure (AW) (Brams and Taylor 1996, 2000), which produces settlements that are efficient, envy-free and equitable with respect to bargainers' *announced* preferences for *n*=2[21]. However, the AW provides no incentives for player to be truthful about their preferences: it is once again the information structure which determines the properties of the solution. In this procedure, two parties begin by independently distributing a total of 100 points across all items to be allocated, according to their own valuation of the goods. Each player is then assigned the goods which s/he values most. The initial allocation is then adjusted to equalise the total valuations of the goods for the two players[22]. The allocation thus achieved is efficient – no player can be made better off without the other being worse off; it is equitable, in that announced valuations are equated; and it is envy-free – no player would trade his or her allocation for that of the other player. However, envy-freeness and equitability are only apparent, as they rely on the truthful revelation of players' valuations – with asymmetric information, the player with complete information can exploit the other player and manipulate the procedure.

---

[21] The AW, in its basic form, implements the Kalai–Smorodinsky solution (Raith 2000).
[22] The same argument applies to issues negotiated (continuous vs. discrete), where player 1 gets 60% of the issue means that the issue is resolved 60% in favour of player 1.

An alternative envy-free allocation – which is however not efficient – is the Proportional Allocation (PA). As the name indicates, under this procedure players are allocated the same share of their valuation of the goods, hence the resulting allocation is envy-free. Under PA, players have the incentive to reveal near true preferences, as the payoffs are hardly affected by deviations[23].

PA's incentives to be truthful come at an efficiency cost with respect to AW: it is however possible to induce players to reveal their nearly-true valuation under an AW procedure, by imposing a PA allocation as a default, should either player be dissatisfied with the allocation reached under AW.

An alternative is Raith's Adjusted Knaster (AK) procedure (Raith, 2000), which is a combination of the Knaster and the AW procedures. AK combines the efficient side payments of the sealed bid procedure with the equitability conditions of AW, for 2 players. By imposing an equitable monetary transfer, the AK implements an outcome that is at least as good as that of the AW.

The fact that these procedures ensure efficiency, equitability and envy-freeness in the 2 persons case is encouraging, despite AW's theoretical (but probably not practical [24]) vulnerability to strategic manipulation. Unfortunately, neither AW nor PA maintain these properties when there are $n>2$ players. Algorithms have been developed that find an allocation satisfying two of the three properties: which pair of properties constitutes the most desirable set is not clear a priori.

When the number of players increases, the algorithms for envy-free allocation get very complicated. There are modifications of other procedures which generate near envy-free allocations, with the degree of error being within any present tolerance level (Brams and Taylor, 1996, p.129). For instance, the general moving knife procedure can be modified to allow players to re-enter the game even though they have received a piece of the cake, by calling cut again and again, with the provision that they must take the piece of cake determined by their most recent cut, and return the previous piece.

---

[23] In fact, it is shown in Brams and Taylor (1996), p. 77, that, in the absence of reliable information about the opponent's preferences, it is a dominant strategy for each player to reveal a valuation close to his or her true valuation – especially in the range 20-80. For extreme values, truth revelation is a dominant strategy if players have symmetric or opposite valuation of the good.
[24] See Brams and Taylor (1996), p. 85.

In a generalisation of the divide-and-choose procedure, players can achieve an envy-free allocation of part of the cake[25] through a trimming strategy: at different stages of the game, parties create equal shares for themselves by trimming others' partitions of the surplus. Note that the final allocation will depend on the order in which players move – there are therefore many envy-free allocations, and many possible equilibria to this game. Moreover, not only does this procedure become very complicated as *n* increase, but also it is not clear what to do with the trimming and the piece left aside: these cannot be distributed, as in the case for *n*=3, in a manner that leaves players envy-free, and exhausts the cake. If one allows for an infinite number of stages, then the procedure can be applied over and over again – and eventually the whole cake is allocated. But the corresponding finite algorithm is complex and unbounded in the sense that the number of cuts needed to produce a given division depends on the number of players *and* on their preferences. Within this procedure, the existence of an envy-free allocation rests on the assumption that the good is divisible – or, in the case of indivisible goods, that there are enough of more divisible goods which can be trimmed in lieu of the discrete good.

Brams and Kilgour (2001) propose a fair division procedure for the allocation of indivisible goods and divisible bads (the price to be paid for the goods). In the Gap procedure, goods are assigned to players in such a way that the total sum of their bids for the goods is maximised (*maxsum* allocation); the prices paid are obtained by decreasing bid values to the next highest bids until their sum is equal to or less than the total value of the goods: once this level is reached (and provided that the sum is not equal to it) reductions in the next higher bids are made in proportion to the *differences* between these bids for each good and the next lower bids. Bids therefore serve the dual purpose of assigning the goods to the players, and determining the prices players have to pay: however, unlike in the Knaster's procedure, the highest bidders do not necessarily receive the goods, and the prices players pay depend not only on their own evaluation of the good, but also on other players' – the more competitive the bids, the higher the price to pay will be. Under the Gap procedure, no players ever pays a negative price – the lowest price a player can pay being the lowest bid for that good; no player pays

---

[25] This is an extension of the Selfridge-Conway discrete procedure, which allows an envy-free allocation of a heterogeneous good among three players with different valuation of the good in question.

more than his bid – players pay either their own bid, or a lower price; the allocation is Pareto efficient, because it maximises total surplus.

However, the Gap procedure does not produce envy-free allocations – that is, a player may prefer the good assigned to another player, at the price that the other player pays for it, to the one assigned to himself, at the Gap price. Potthoff (2002) proposes a linear programming solution to find an envy free solution which is closes to the Gap solution – that is, that set of envy-free prices that minimises the sum of absolute difference from the Gap prices. Such a solution always exists when negative prices are allowed – but its existence is not guaranteed otherwise

6.4. Synthesis of the procedures

Whatever the underlying motivations are for the emergence of "fair" behaviours and/or outcomes, the *perception* of fairness is critical to facilitate the achievement of an agreement on how to divide a surplus in a non-cooperative negotiation framework, where allocations need to be self-enforcing.

In

Table 2, the main fair division procedures presented here are summarised and compared, with respect to three main characteristics: equity, envy-freeness, and efficiency. An allocation is equitable when players think that their portion is worth the same as everybody else's'; it is envy-free when every player thinks s/he receives a portion that is at least tied for the largest, or tied for most valuable, and hence does not envy any other player; and it is efficient, if no player can be made better off, without another player being made worse off.

When the properties of the fair division procedures vary depending on the number of players (two or more than two) and/or the type of item they can be applied to (homogenous or heterogeneous, divisible or indivisible), this is emphasised in the table.

Table 2: Summary of the main fair division procedures, and key characteristics

	Players	Surplus	Equity	Envy-free	Efficient
Basic Fair Division Procedures					
Discrete procedures					

Basic divide and choose	2	Divisible Heterogeneous	Proportional	Yes – if no asymmetries in information	No unless players have symmetric preferences for all the parts of the cake
Filer and choose	2	Non-divisible – public good	Proportional	Yes – if no asymmetries in information	No unless players have symmetric preferences for all the parts of the cake
Discrete trimming	>2	Non-divisible	Proportional	Yes	No
Selfridge-Conway discrete	3	Divisible, heterogeneous	Proportional	Yes	No
Lone-divider	>2	Divisible Heterogeneous	Proportional	No	No unless players have symmetric preferences for all the parts of the cake
Lone-chooser	2 and >2	Divisible	Proportional	Yes for 2 players. Otherwise no.	No unless players have symmetric preferences for all the parts of the cake
Continuous procedures					
Moving knife	2	Divisible Heterogeneous	Proportional	Yes – if no asymmetries in information	No unless players have symmetric preferences for all the parts of the cake

Generalised moving knife	>2	Divisible	Proportional	No	No unless players have symmetric preferences for all the parts of the cake
Last diminisher	>2	Divisible	Proportional	No	No unless players have symmetric preferences for all the parts of the cake
Refinements of the basic procedures					
Adjusted winner	2	Divisible Indivisible	Proportional	Yes – with respect to players *stated* preferences	Yes
	>2	Divisible Indivisible	It can satisfy two of the properties only		
Proportional allocation	2 and >2	Divisible	Proportional	Envy-free	No
Adjusted Knaster's procedure	>2	Divisible Indivisible	Proportional	Envy-free	No
Gap procedure (max-sum allocation)	>2	Indivisible goods and divisible bads	Proportional	Envy-free	Yes

The procedures described in this short review are applicable to both homogenous and heterogeneous players – it is in fact the structure of information which determines the properties of the solution. When (a)symmetry of information and players' preference structure affect the properties of the solution, this is highlighted in the table.

There are therefore numerous fair-division procedures, which exhibit different properties with respect to the efficiency, equitability, envy-freeness of both the procedures and the resulting allocation. It is difficult to answer theoretically which procedure is best, as trade-offs among their characteristics, as

well as consideration of vulnerability of the procedure to strategic manipulation, need to be considered.

However, the focus of researcher and practitioners should shift away from the achievement of an efficient allocation as the overriding goal, and pay more attention to the properties of equity and envy-freeness – which should be satisfied, if a self-enforcing agreement is needed. In fact, restricting the possible agreements to those satisfying some form of equity and envy-freeness could help select one equilibrium when a multiplicity of equilibria could be possible.

7   Conclusions

The relevance of negotiations to everyday life cannot be overemphasised. Yet, a comprehensive theory of negotiation is still missing: the factors involved in the processes of negotiations are so complex and varied, that they have been tackled in isolation, with the consequence that many theoretical results of the standard models do not always find support in empirical evidence.

From this review of the theory four main considerations emerge, which should be taken into account in the formulation of a suitable negotiation model:

- The *non-cooperative approach* to negotiations is useful in that it allows for the analysis of players' incentives to cooperate. Moreover, the outcome of a non-cooperative game has the property of being self-enforcing. This is particularly important at an international level where there are no supranational governing bodies which can impose cooperation, and agreements have to be reached voluntarily among sovereign states.

- The *sequential-move approach* enables the process of negotiation to be modelled. This, in turn, allows for the analysis of some particular issues (such as bargaining and political power, asymmetric information, time preferences) which may have relevant effects on the bargaining outcome.

- However, standard bargaining theory is not well suited to deal with bargaining situations where (positive) *externalities* are involved. The presence of externalities opens up the possibility of intermediate agreements, neither fully cooperative, nor fully non-cooperative. These more complex situations can be better explored by non-cooperative coalition theory.

- Finally, both standard bargaining theory and coalition theory do not address the issue of *fair division* in a comprehensive manner, focusing almost exclusively

on the efficiency property of the outcomes. The integration of fair division theory in negotiation is however crucial if the solution/agreement is to be implemented and sustained.

8   References

1. Alesina, A., I. Angeloni and F. Etro (2001), 'The political economy of unions', NBER Working Paper, December 2001.

2. Andreoni, J. and J. H. Miller (1996), 'Giving according to GARP: an experimental study of rationality and altruism', *SSRI Working Paper, University of Wisconsin*.

3. Bac, M. and H. Raff (1996), 'Issue-by-issue Negotiations: the Role of Information and Time Preferences, *Games and Economic Behavior*, 13, 125-134.

4. Barrett, S. (1994), 'Self-Enforcing international Environmental Agreements', *Oxford economic Papers*, 46, 804-878.

5. Barrett, S. (1995), 'Trade restrictions in international environmental agreements', London Business School.

6. Barrett, S. (1997), 'Heterogeneous International Agreements', in C. Carraro (ed.), "International Environmental Negotiations: Strategic Policy Issues", Edward Elgar, Cheltenham, pp. 9-25.

7. Beersma, B. and C. K. W. de Dreu (2003), 'Social motives in integrative negotiation: the mediating influence of procedural fairness', *Social Justice Research* 16, 217-239.

8. Bloch, F. (1996), 'Sequential Formation of Coalitions in Games with Externalities and Fixed Payoff Division', *Games and Economic Behaviour*, 14, 90-123.

9. Bloch, F. (1997), 'Non-cooperative Models of Coalition Formation in Games with Spillovers', in Carraro, C. and D. Siniscalco (eds.), "New Direction in the Economic Theory of the Environment", Cambridge University Press, Cambridge, Ch. 10, pp. 311-352.

10. Bloch, F. and A. Gomes (2003), 'Contracting with Externalities and Outside Options', discussion paper at the *IX Coalition Theory Network*, Barcelona, 30th -31st January, 2004.

11. Brams, S. J. and M. D. Kilgour (2001), 'Competitive fair division', *Journal of Political Economy* 109, 418-443.

12. Brams, S. J. and A. D. Taylor (1995), 'An envy-free cake division protocol', *American Mathematical Monthly* 102, 9-18.

13. Brams, S. J. and A. D. Taylor (1996), 'Fair division: from cake cutting to dispute resolution', Cambridge University Press, Cambridge, UK.

14. Brams, S. J. and A. D. Taylor (2000), 'The Win-Win solution: guaranteeing fair shares to everybody', W.W. Norton & Company, New York.

15. Busch, L. A. and I. J. Horstmann (1997), 'A Comment on Issue-by-Issue Negotiations', *Games and Economic Behavior*, 19, 144-148.

16. Camerer, C. and R. Thaler (1995), 'Ultimatum, dictators and manners', *Journal of Economic Perspectives* IX, 209-219.

17. Carraro, C. and D. Siniscalco (1993), 'Strategies for the International Protection of the Environment', *Journal of Public Economics*, 52, 309-328.

18. Carraro, C. and D. Siniscalco (1995), 'Policy coordination for sustainability: commitments, transfers and linked negotiations', in I. Goldin and A. Winters (eds), *The Economics of Sustainable Development*, Cambridge: Cambridge University Press.

19. Carraro, C. and D. Siniscalco (1997), 'R&D cooperation and the stability of international environmental agreements', in C. Carraro (ed), *International Environmental Agreements: Strategic Policy Issues*, Cheltenham: Edward Elgar.

20. Cesar, H. and A. De Zeeuw (1996), 'Issue linkage in global environmental problems', in A. Xepapadeas (ed), *Economic Policy for the Environment and Natural Resources*, Cheltenham: Edward Elgar.

21. Chae, S. and J. A. Yang (1988), 'The Unique Perfect Equilibrium of an *N*-person Bargaining Game', *Economic Letters*, 28, 221-223.

22. Chae, S. and J. A. Yang (1994), 'An *N*-person Pure Bargaining Game', *Journal of Economic Theory*, 62, 86-102.

23. Chander, P. and H. Tulkens (1997), 'The Core of an Economy with Multilateral environmental Externalities', *International Journal of Game Theory*, 26, 379-401.

24. Chatterjee, K., B. Dutta, D. Ray and K. Sengupta (1993), 'A Non-cooperative Theory of Coalitional Bargaining', *Review of Economic Studies*, 60, 463-77.

25. D' Aspremont, C., A. Jaquemin, J.J. Gabszewic and J. Weymark (1983), 'On the Stability of Collusive Price leadership', *Canadian Journal of Economics*, 16, 17-25.

26. Fatima, S., M. Wooldridge and N. R. Jennings (2002), 'Bargaining with Incomplete Information, mimeo, University of Southampton.

27. Fatima, S., M. Wooldrige and N. R. Jennings, 'Optimal Agendas for Multi-Issue Negotiation', mimeo, University of Southampton.

28. Fehr, E. and K. M. Schmidt (1999), 'A theory of fairness, competition and cooperation', *The Quarterly Journal of Economics* 114, 817-849.

29. Fershtman, C. (1990), 'The Importance of the Agenda in Bargaining', *Games and Economic Behavior*, 2, 224-238.

30. Folmer, H., P. van Mouche and S. Ragland (1993), 'Interconnected games and

international environmental problems', *Environmental Resource Economics*, 3, 313-35.

31. Forsythe, R., N. E. Horowitz, H. L. Savin and M. Sefton (1988), 'Fairness in simple bargaining games', *Games and Economic Behaviour* 6, 347-369.

32. Fudenberg, D. and J. Tirole (1983), 'Sequential Bargaining with Incomplete Information', *Review of Economic Studies*, 50, 221-247.

33. Fudenberg, D. and J. Tirole (1985), 'Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information', in A. Roth (ed.), *Game Theoretic Models of Bargaining*, University of Cambridge Press, Cambridge.

34. Gomes, A. (2001), 'Multilateral Contracting with Externalities', mimeo, The Wharton School, University of Pennsylvania.

35. Güth, W., R. Schmittberger and R. Tietz (1990), 'Ultimatum bargaining behaviour: a survey and comparison of experimental results', *Journal of Economic Psychology* XI, 417-449.

36. Harrison, G. and K. McCabe, 'Expectations and fairness in a simple bargaining experiment', New Orleans LA. Paper presented at The Public Choice Society Meeting.

37. Harsanyi, J. C. And R. Selten (1972), 'A Generalized Nash Solution for two-Person Bargaining Games with Incomplete Information', *Management Science*, 18, 80-106.

38. Hart, S. and M. Kurz (1983), 'Endogenous Formation of Coalitions', *Econometrica*, 51, 1047-64.

39. Herrero, M. (1985), 'A Strategic Theory of Market Institutions', unpublished doctoral dissertation, London School of Economics.

40. Hoel, M. (1992), 'International Environmental Conventions: The Case of Uniform Reductions of Emissions', *Environmental and Resource Economics*, 2, 141-159.

41. Hoel, M. and K. Schneider (1997), 'Incentives to Participate in an International Environmental Agreement', *Environmental and Resource economics*, 9, 153-170.

42. In and Serrano (2003), 'Agenda Restrictions in Multi-Issue Bargaining: Unrestricted Agendas', *Economics Letters*, 79, 325-331.

43. Inderst, R. (2000), 'Multi-issue Bargaining with Endogenous Agenda', *Games and Economic Behavior*, 30, 64-82.

44. Jun, B. H. (1987), 'A Structural Consideration on 3-person Bargaining', *Chapter III in Essays on Topics in economic Theory,* Ph. D Thesis, Department of Economics, University of Pennsylvania.

45. Katsoulacos, Y. (1997), 'R&D spillovers, R&D cooperation, innovation and

international environmental agreements', in C. Carraro (ed), *International Environmental Agreements: Strategic Policy Issues*, Cheltenham: Edward Elgar.

46. Kravitz, D. and S. Gunto (1992), 'Decisions and recipients in ultimatum bargaining games', *Journal of Social Economics* 21, 65-84.

47. Krishna, V. and R. Serrano (1996), 'A Model of Multilateral Bargaining', *Review of Economic Studies,* 63, 61-80.

48. Loewenstein, G. F., L. Thompson and M. H. Bazerman (1989), 'Social utility and decision making in interpersonal contexts', *Journal of Personality and Social Psychology* 57, 426-441.

49. Marteens, M., J. Potters and H. Reijnierse (2002), 'Envy-free and Pareto efficient allocations in economies with indivisible goods and money', *Mathematical Social Sciences* 44, 223-233.

50. Moldovanu, B. (1992), 'Coalition-Proof Nash Equilibria and the core in Three-Player Games', *Games and Economic Behaviour*, 4, 565-81.

51. Montero, M. (1999), 'Coalition Formation in Games with Externalities', *CentER Discussion Paper*, 99121.

52. Muthoo, A. (1995), 'Bargaining in a Long-Term Relationship with Endogenous Termination', *Journal of Economic Theory*, 66, 590-598.

53. Ochs, J. and A. E. Roth (1989), 'An experimental study of sequential bargaining', *The American Economic Review* 79, 355-384.

54. Okada, A. (1996), 'A Non-cooperative Coalitional Bargaining Game with Random Proposers', *Games and Economic Behaviour*, 16, 97-108.

55. Perry, M. and P. Reny (1994), 'A Non-cooperative View of Coalition Formation and the Core', *Econometrica*, 62, 795-817.

56. Petrakis, E. and A. Xepapadeas (1996), 'Environmental consciousness and moral hazard in international agreements to protect the environment', *Journal of Public Economics*, 60, 95-110.

57. Pillutla, M. M. and J. K. Murnighan (2003), 'Fairness in bargaining', *Social Justice Research* 16, 241-262.

58. Potthoff, R. F. (2002), 'Use of linear programming to find an envy-free solution closest to the Brams-Kilgour Gap solution for the housemates problem', *Group Decision and Negotiation* 11, 405-414.

59. Rabin, M. (1993), 'Incorporating fairness into game theory and economics', *American Economic Review* 83, 1281-1302.

60. Raith, M. G. (2000), 'Fair negotiation procedures', *Mathematical Social Science* 39, 303-32.

61. Ray, D. and R. Vohra (1996), 'Binding Agreements and Coalitional Bargaining', mimeo, Department of Economics, Brown University and Boston University.

62. Ray, D. and R. Vohra (1999), ' A Theory of Endogenous Coalition Structures', *Games and Economic Behaviour*, 26, 286-336.

63. Roth, A. E. (1995), 'Bargaining experiments', in J. Kagel and A. E. Roth (eds.), "Handbook of Experimental Economics", Princeton University Press, Princeton.

64. Rubinstein, A. (1985), 'A Bargaining Model with Incomplete Information about Time Preferences', *Econometrica*, 53, 1151-1172.

65. Rubio, S. and A. Ulph (2001), 'A Simple Dynamic Model of International Environmental Agreements with a Stock Pollutant', mimeo, University of Southampton.

66. Sandholm, T. and N. Vulkan, (1999), 'Bargaining with Deadlines', in *AAAI-99*, Orlando, FL, 44-51.

67. Sebenius, J.K. (1983), 'Negotiation arithmetic: adding and subtracting issues and parties', *International Organisation*, 37, 281-316.

68. Seidmann, D. and E. Winter (1998), 'A Theory of Gradual Coalition Formation', *Review of Economic Studies*, 65, 793-815.

69. Selten, R. (1981), 'A Non-cooperative Model of Characteristic Function Bargaining', in V. Bohm and H. Nachtkamp (eds.), "Essays in game Theory and mathematical economics in Honor of Oskar Morgenstern", Mannheim: Bibliographisches Institut Mannheim, pp. 131-151

70. Shaked, A. (1986), 'A Three-Person Unanimity Game', mimeo.

71. Slonim, R. and A. E. Roth (1997), 'Financial incentives and learning in ultimatum and market games: an experiment in the Slovak Republic', *Econometrica* 66, 569-596.

72. Straub, P. and J. K. Murnigham (1995), 'An experimental investigation of ultimatum: complete information, fairness, expectations, and lowest acceptable offers', *Journal of Economic Behaviour and Organisations* 27, 345-364.

73. Thaler, R. H. (1988), 'The ultimatum game', *Journal of Economic Perspective* II, 195-206.

74. Tollison, R.D. and T.D. Willett (1979), 'An economic theory of mutually advantageous issue linkages in international negotiations', *International Organisation*, 33(4), 425-49.

65. Yi, S. (1997), 'Stable coalition structures with externalities', *Games and Economic Behaviour*, 20, 201-23.

66. Yi, S. And H. S. Shin (1994), 'Endogenous formation of coalitions in oligopoly', Dartmouth College, Working Paper 94-18