

DISCUSSION PAPER SERIES

No. 6424

**MODELLING NEGOTIATED
DECISION MAKING: A
MULTILATERAL, MULTIPLE ISSUES,
NON-COOPERATIVE BARGAINING
MODEL WITH UNCERTAINTY**

Carlo Carraro and Alessandra Sgobbi



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP6424.asp

MODELLING NEGOTIATED DECISION MAKING: A MULTILATERAL, MULTIPLE ISSUES, NON-COOPERATIVE BARGAINING MODEL WITH UNCERTAINTY

Carlo Carraro, University of Venice, Fondazione Eni Enrico Mattei,
EuroMediterranean Center on Climate Change and CEPR
Alessandra Sgobbi, Fondazione Eni Enrico Mattei and EuroMediterranean Center
on Climate Change

Discussion Paper No. 6424
August 2007

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in . Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Carlo Carraro and Alessandra Sgobbi

ABSTRACT

Modelling Negotiated Decision Making: A Multilateral, Multiple Issues, Non-Cooperative Bargaining Model with Uncertainty*

The relevance of bargaining to everyday life can easily be ascertained, yet the study of any bargaining process is extremely hard, involving a multiplicity of questions and complex issues. The objective of this paper is to provide new insights on some dimensions of the bargaining process – asymmetries and uncertainties in particular – by using a non-cooperative game theory approach. We develop a computational model which simulates the process of negotiation among more than two players, who bargain over the sharing of more than one pie. Through numerically simulating several multiple issues negotiation games among multiple players, we identify the main features of players' optimal strategies and equilibrium agreements. As in most economic situations, uncertainty crucially affects also bargaining processes. Therefore, in our analysis, we introduce uncertainty over the size of the pies to be shared and assess the impacts on players' strategic behaviour. Our results confirm that uncertainty crucially affects players' behaviour and modify the likelihood of a self-enforcing agreement to emerge. The model proposed here can have several applications, in particular in the field of natural resource management, where conflicts over how to share a resource of a finite size are increasing.

JEL Classification: C61, C71 and C78

Keywords: bargaining, non-cooperative game theory, simulation models and uncertainty

Carlo Carraro
Dipartimento di Scienze Economiche
Università di Venezia
San Giobbe 873
30121 Venezia
ITALY
Email: ccarraro@unive.it

Alessandra Sgobbi
FEEM
Corso Magenta 63
20121 Milan
Italy
Email: alessandra.sgobbi@feem.it

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=104121

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=166881

Submitted 30 July 2007

Introduction

The study of any bargaining process is extremely hard, involving a multiplicity of questions and complex issues. Crucial questions to be addressed are, for example: What factors determine the outcome of negotiations? What strategies can help reach an agreement? How should the parties involved divide the gains from cooperation? With whom will one make alliances? The research literature in this field has not yet succeeded in developing a comprehensive framework for analysis, and a number of theories have been proposed instead, each focusing on single aspects of the problem. As a consequence, theoretical results are not always supported by empirical evidence.

The objective of this paper is to provide new insights on some dimensions of the bargaining process – asymmetries and uncertainties in particular – by using a non-cooperative game theory approach. We develop a computational model which simulates the process of negotiation among more than two players, who bargain over the sharing of more than one pie. We then explore the effects that uncertainty over the size of the pies has on players' strategies and equilibrium agreements.

The remainder of this paper is organised as follows. Section 2 provides a brief survey of the related literature, describing the main progresses and the research gaps that still need further efforts. Section 3 introduces the underlying bargaining framework, while the results of numerically solving our dynamic bargaining model are summarised in Section 4. Sensitivity analysis is carried out in Section 5 whereas, in Section 6, we introduce uncertainty over the negotiated variables, and explore the implications for players' strategic behaviour and equilibrium agreements. Section 7 concludes the paper.

Brief survey of the literature

The basic bargaining model introduced by Rubinstein (1982) describes the process through which negotiating agents try to reach an agreement as a situation in which players make offers and counter-offers over the terms of the agreement. Rubinstein's model is an infinite horizon model of perfect information, which extends the work of Ståhl (1972) – a finite horizon alternating offers model – and takes explicitly into account players' strategic incentives to cooperate. In the game, two players take turn in making a proposal over the sharing of a resource of known size: if players reach an agreement over the allocation, the game ends; if not, the game moves on to the next stage, when it is the other player's turn to make a proposal. The model captures an intuitive process of bargaining and, although extremely simplified, provides the formal basis for many important results. In particular, it clarifies how different individuals' characteristics affect the bargaining outcome.

Rubinstein proves that every bargaining game of alternating offers that satisfies the assumptions of his model has a unique sub-game perfect equilibrium (SPE)¹, which is reached immediately. From an economic point of view, the fact that negotiation ends in the first period implies that the equilibrium is efficient because no resources are lost in delay. Furthermore, the model predicts that when a player becomes more patient (that is, when he values the future more) relative to the opponent, his bargaining power increases, and so does his share of the cake². Thus, in this basic formulation, the bargaining power depends on players' relative cost of waiting³.

In many real life situations, however, negotiations involve more than two players, and may involve more than one issue. Therefore, in recent years, several extensions of the standard Rubinstein's model have been proposed in the attempt to deal with more complex bargaining situations and to find theoretical explanations for the observed empirical regularities.

2.1 Multiple players

When there are more than two players at the negotiation table, the characteristics and equilibrium solutions predicted by the basic Rubinstein model do not hold anymore. Even in the case of three players, the most natural extension of the Rubinstein bargaining structure leads to indeterminate results. The indeterminacy of the $n \geq 3$ player game has aroused much interest among researchers and various solutions have been proposed to arrive at a unique equilibrium outcome – either by adopting different (more refined) equilibrium concepts, or by modifying the structure of the game.

For instance, by limiting the strategies available to players to stationary strategies (Shaked, 1986), it is possible to obtain an equilibrium similar to the unique SPE of the two-player game.⁴ The notion of stationary SPE may therefore be used to restore the uniqueness of the equilibrium in multilateral bargaining situations, but the restriction on the strategy space is rather strong. A more appealing way to address the problem of indeterminate result with $n \geq 3$ is to modify the structure of the game. For example, Jun (1987) and Chae and Yang (1988, 1994) propose a process where players are engaged in a series of bilateral negotiations, and any player that reaches a satisfactory agreement

¹ A strategy profile is a Sub-game Perfect Equilibrium of a game if it is a Nash equilibrium of every sub-game of the game. More formally, for any history of the game G , h^k , a strategy s is sub-game perfect if, for any restriction $s|h^k$, it is an equilibrium of the game $G(h^k)$.

² It is important to note that, when players do not discount payoffs, then any partition of the cake can be agreed upon at time 0, and can be supported in an SPE. In this case, neither players care about the time at which the agreement is struck – that is, negotiation is frictionless.

³ As we shall see later on, players' bargaining power may be influenced by many other factors, including the utility they derive from not reaching an agreement.

⁴ Stationary strategies prescribe actions in every period that do not depend on time, nor on events in previous periods – that is, strategies are history independent. Herrero (1985) shows that if players have time preferences with a common, constant, discount factor, there is a unique allocation of the pie among the three players, which tends to an equal split as players become more patient.

may “exit” the game. In Krishna and Serrano (1996) players have the possibility to leave the negotiations with their share before all players reach an agreement, but all offers are made to all players simultaneously – that is, bargaining is multilateral. With the introduction of such rules, the authors are able to identify a unique perfect equilibrium for any number of players.⁵

More recently, Manzini and Mariotti (2005) examine a different set of bargaining rules. Their model is designed as a two-stage game: the “outer” stage is the standard Rubinstein model, while in the “inner stage” members of a coalition have to agree among themselves over the choice of proposal to offer to the other party. The main result of their work is that unanimity in the inner stage seems to lead coalitions to adopt a less flexible stand over their opponent than majority rule. In another work (Manzini and Mariotti, 2003), the authors find that, when unanimity is required, the outcome of the negotiation is entirely determined by the toughest negotiator.

2.2 Multiple issues

Many real-life negotiations do involve a set of different issues, but most of the existing literature on non-cooperative games focuses on single issue negotiations. When all issues are bundled together in a negotiation package (complete package approach), standard theory still applies. Alternatively, multiple issues can be negotiated one by one (sequential approach). 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 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. The author shows that a player prefers to bargain first over the least important project to him, if it is the most important for his opponent. Furthermore, as players become more patient (and the cost of bargaining decreases) the impact of the agenda vanishes.

The agenda is partially endogenised in Busch and Horstmann (1997), where the bargaining game is preceded by a separate bargaining round over the order of projects’ negotiation. In the work by Inderst (2000), the agenda of accepted projects becomes truly endogenous: players are allowed to freely choose the subset of projects for which they want to make an offer. The analysis reveals that the agenda can have a significant impact on payoffs and this impact does not seem to vanish as players become increasingly patient. It emerges that (i) bargaining simultaneously over a set of projects can improve efficiency by creating trading opportunities across issues; (ii) changing the agenda may have a distributive effect, and players may therefore prefer different agendas. Bac and Raff (1996) investigate how incomplete information about bargaining strength impacts the choice of the bargaining

⁵ The original model proposed by Krishna and Serrano (1996) has been extended to include the existence of interpersonal externalities by Kultti and Vartiainen (2004).

procedure, where players' bargaining strength is given by their time preferences, and players negotiate over more than one issue. The model involves two players negotiating à la Rubinstein over the division of two identical pies of size 1, and parties' offers include the bargaining procedure (issue by issue vs. simultaneous), as well as an allocation. The model proposed has an asymmetric information structure, where one player is perfectly informed about the time preferences of both, whereas the other player is uncertain about the time preference of his opponent. With two possible realisations of the (uncertain) discount factor, the bargaining game has a sequential equilibrium with rationalising beliefs such that, while a weak player (high discount rate) prefers to negotiate simultaneously, a strong player (low discount rate) bargains sequentially to signal his bargaining strength.

In a more recent work, In and Serrano (2003) develop a model to investigate 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 delays in agreement do usually arise.

Lastly, Fatima et al. (2003) study the strategic behaviour of agents by using an agenda that is partly exogenous and partly endogenous, in an attempt to determine whether agents' utilities can be improved by decomposing the set of issues into stages, allowing exogenous determination of issues for each stage, and endogenous, sequential negotiation of each stage. The analysis shows that the optimal number of decompositions for an agent depends on his negotiation parameters. In other words, there exist negotiation scenarios where the utility of 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 previous works, namely that differing preferences over issues play an important role in determining negotiation agendas.

2.3 Uncertainty

Most of the traditional models of bargaining deal with issues known with certainty. In many real-life situations, however, the issues negotiated over may not be certain, but follow a stochastic process. Even though each player is perfectly informed about his opponent's parameters, there is a source of uncertainty that can affect the negotiation outcome. In general, exogenous uncertainty may enter into a bargaining process in many different ways. For example, the size of the surplus over which players are negotiating may vary stochastically, as well as the disagreement point; the identity of the proposing player may also follow a random process; unexpected breakdowns or unexpected changes in bargaining positions may occur, and so on. Let us focus on the case in which uncertainty affects the size of the pie, as this is the case that will be analysed in the subsequent sections.

Merlo and Wilson (1995)⁶ propose an extension of the basic Rubinstein two-players 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 are stochastic, exogenously determined, and follow a general Markov process. Using the concept of stationary sub-game perfect (SSP) equilibria⁷, the authors find that there exist a unique SSP equilibrium, which is efficient, even though it may involve (efficient) delays. This result does not exactly conform 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, or outcomes with delay may arise but efficiency is not guaranteed anymore. In this model, agreement may be delayed even in the unique SSP 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 because 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.

Eraslan and Merlo (2002) propose an extension of Merlo and Wilson (1995) to allow for general agreement rules, studying a multilateral stochastic bargaining game of complete information with q -quota agreement. It has been shown that bargaining games with general q -quota agreement rule but fixed surplus (deterministic environment) usually admit a unique stationary sub-game perfect equilibrium which is efficient and involves no delays. This result, however, is not maintained in games with uncertainty over the size of the payoff. In particular, Eraslan and Merlo (2002) find that the uniqueness of the equilibrium is not guaranteed anymore and, even when the equilibrium is unique, for any agreement rule other than unanimity it need not be efficient. The kind of inefficiency that may emerge at the equilibrium is induced by the fact that agreement may be reached “too soon”. Intuitively, whenever an agreement entails less than unanimous approval, there exists a differential treatment between the players who are included in a proposal and the players who are excluded. In a stochastic environment there may be incentives for the players to delay agreement until a larger level of surplus is realised. This tension generates the possibility of inefficient agreements where players fail to realise all the gains from waiting, and may also generate multiplicity of equilibrium payoffs. Players who are offered a positive payoff in a state where the level of surplus is relatively small may be induced to accept it if they expect to be excluded from future agreements when the level of surplus is relatively large.

⁶ Various applications exist of this framework, which 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.

⁷ Recall that a strategy profile is *sub-game perfect* (SP) if, at every history, it is a best response to itself, and is *stationary* (SSP) if the actions prescribed at any history depend only on the *current* state and *current* offer.

What are then the key issues that should be considered in designing a model to simulate negotiation processes among several players over several issues, and extending it to include uncertainty over the negotiated variables?

First of all, when multiple players are involved in a non-cooperative negotiation process, the standard result of a unique equilibrium agreement do not hold any longer. To restore uniqueness, thus improving the usefulness of the models in exploring negotiation strategies and predicting possible outcomes, more stringent negotiation rules need to be applied, or, alternatively, one needs to impose refinements of the equilibrium outcomes. Secondly, when players bargain over more than one policy dimension, simultaneous bargaining is to be preferred: in line with the theory of issue linkage (see, for instance, Folmer et al., 1993; Cesar and De Zeeuw, 1996; Carraro and Siniscalco, 1997; and, more recently, Alesina et al, 2001), simultaneous bargaining enlarges the zone of agreement, thus creating gains from trade and increasing the likelihood of an equilibrium to the negotiation process. Finally, when one considers the possibility of stochastic realisations of one or several of the negotiated variables, convergence of players' strategies to a unique (efficient) equilibrium is more difficult: to improve chances of an agreement to emerge, therefore, more stringent decision rules can be imposed, such as unanimity.

The negotiation framework

The above three factors have been taken into account in developing a non-cooperative bargaining model that could usefully simulate the negotiation process and strategic interactions among several players negotiating simultaneously over several issues, when the size of the resource to be allocated may not be known with certainty. Players' strategies will then depend on the expected realisation of future states of the world. If players fail to reach an agreement by an exogenously specified deadline, a disagreement policy is imposed. The disagreement policy is known to all players: it can either be an allocation that is enforced by managing authorities; it can be the loss of the possibility to enjoy even part of the negotiated variable; or it could be the continuation of the status quo, which is often characterised as inefficient and, thus, to be improved.

The constitution of the game as a finite horizon negotiation is justifiable empirically – as consultations over which policies to implement cannot continue forever, but policy makers have the power (if not the interests) to override stakeholders' positions and impose a policy, if negotiators fail to agree. This is also true in negotiation when an arbitrator is involved, who, with the agreement of all parties, has the power to take a final decision regarding the negotiated variables, should agents fail to reach an agreement. In finite horizon strategic negotiation models, it is unavoidable that “11th hour” effects play an important role in determining the equilibrium solution. In fact, last minute agreements are often reported in negotiation settings – think, for instance, about labour agreements which tend to

be reached just before the contract expires, or just before the set day for strikes. In some settings, the possible explanation of the “11th hour effect” appears straightforward. When the profit to be divided does not decrease over time, the outcome predicted by game theory is determined by the anticipated behaviour of players in the last period: the last proposer may hope to ripe (almost) all the benefits by making an ultimatum offer (Ma and Manove, 1993). Starting from the final round of the game and working back to the first period reveals that the subgame perfect equilibrium gives virtually all the pie to the last proposer. Under these conditions, the strategic behaviour of players may aim at creating the conditions for issuing an ultimatum. That is, as the deadline is approached, players may have to make larger concessions (Gneezy et al, 2003).

Our model has as a starting point the negotiation framework proposed by Rausser and Simon (1992), but will include uncertainty in the negotiation space and explore how the strategies of players and the emerging equilibrium agreement vary with respect to the deterministic bargaining game. As the model has no closed form solution, we first validate the results of the Rausser and Simon model and its applications (Adams et al., 1996; Simon et al., 2003, 2006; Thoyer et al, 2001), and then explore numerically the impacts of uncertainty over the realisation of a negotiated variable on the equilibrium outcome of the game.

3.1 The multilateral, multiple issues negotiation framework

The constitutional features of our model are as follows: The set of players is I , with typical element i , and the set of feasible policy dimensions is X , with typical element $k \in X$. Players have well-specified utility functions, which compute players’ payoffs for each of the selected policy vectors $\mathbf{x} \in X$. Crucially, players’ utilities will depend on a stochastic variable, \tilde{k} . Let $U_i(\mathbf{x})$ denote this utility functions, which satisfy four key assumptions:

Assumption 1 (A1): X is a convex, compact subset on the K -dimensional Euclidean space, where K denotes the number of issues to be negotiated simultaneously.

Assumption 2 (A2): players’ utility functions are assumed continuous and strictly concave on X , and to satisfy the Von-Neumann-Morgenstern axioms. The assumption of strict concavity implies that players are risk averse – that is, players are reluctant to accept a bargain with an uncertain payoff rather than another bargain with a more certain but possibly lower payoff.

Assumption 3 (A3): players are assumed to have sufficiently different preferences, that is, for each player $i \neq j$, the maximisers of $U_i(\mathbf{x})$ and $U_j(\mathbf{x})$ are distinct. This assumption is invoked to avoid degenerate outcomes. Finally,

Assumption 4 (A4): there exist a policy vector $\mathbf{x} \in X$ such that $U_i(\mathbf{x}) > U_i(\underline{\mathbf{x}}) = U_0$, where $\underline{\mathbf{x}} \in X$ is the distinguished policy vector that is enforced on players, should they fail to reach an

agreement by the exogenously specified terminal time T , and U_0 is the utility of this disagreement policy. This assumption avoids degenerate case in which no agreement is preferred by all players to agreement.

Players negotiate a complete package agreement, which will remain valid for g periods of time – after which they may renegotiate the policy package agreed upon. In this framework, therefore, there is no learning, as uncertainty is not resolved in the course of the negotiation.

Players are selected to submit an offer in an exogenously specified ordered, determined by a vector of access probabilities, ω . A player i , when submitting an offer, can only propose a policy package that belongs to the policy space X , that is, he can only propose feasible policies.

The only admissible coalition is the grand coalition, that is, unanimity is required to reach an agreement. As highlighted in Section 0, unanimity rules increase the chances of a self-enforcing agreement to emerge. Although this may seem excessively restrictive – in some cases, such as government formation, simple or qualified majority rules may be more realistic – unanimity is justifiable empirically when no cooperation is the status quo, when there is no possibility of binding agreements, or enforcement of an agreement is problematic – all cases in which the agreement must be self-enforcing and voluntary. Unanimity may also be appropriate when a compromise among different perspectives is sought.

The game is played as follows. At each round $t < T$, provided no agreement has yet been reached, the player specified by the sequence ω proposes a policy package $\mathbf{x} \in X$. In particular, we will assume that players, when selected to be proposers, maximise their utility by requesting for themselves a share $x_{k,i}$, and proposing to other players an allocation that is minimal with respect to their participation constraint. Next, all the remaining players respond to the offer in the order specified by ω . If all players accept the proposal, the game ends. If there is at least one player that rejects the offer, the next period of the game starts. In $t+1$, the next player in the sequence specified by ω proposes a policy package $\mathbf{x}' \in X$, $\mathbf{x}' \neq \mathbf{x}$, which the remaining players can in turn either accept or reject. The game continues in this fashion until either all players agree to a proposed policy package, or the terminal time T is reached, at which point the disagreement policy $\underline{\mathbf{x}}$ is implemented.

The game is solved through a series of single-person optimisation problems, in which players attempt to maximise their gain from the final agreement over how to share the resources, subject to the total quantity of resources available, and to the agreement being accepted by the other players.

The equilibrium concept for this game is sub-game perfection. As any policy which is weakly preferred by players to the disagreement outcome can be sustained as an outcome, an equilibrium refinement concept is invoked to restore uniqueness of the equilibrium solution. An equilibrium refinement provides a way of selecting one or a few equilibria from among many in a game. Each refinement attempts to define some equilibria as "more likely," "more rational" or "more robust" to

deviations by players than others. For example, if one equilibrium results in all players earning more than another, it may be more likely that the players will coordinate or be naturally drawn to it.

Following the original work by Rausser and Simon (1992), we apply the Sequential Elimination of Dominated Strategies (SEDS) (Myerson, 1978), which eliminates strategies that involve inadmissible (i.e., weakly dominated) play, starting from the final response round, and reaching the first round of negotiation. The outcome generated by such strategy profile is the equilibrium outcome for the game.

3.2 Uncertain surplus

Let us now introduce an element of uncertainty in the game, by assuming that one of the policy issues negotiated over varies stochastically, and its realisation is not known with certainty.

As we are interested in allocation agreements which have a specified duration, the standard assumption in stochastic programming (with recurse) that uncertainty will be resolved in a second stage of the game cannot hold. In stochastic programming, players are assumed to take decisions accounting for uncertainty, that is, by assuming a probability distribution over possible realisations of different states of the world. Players are thus assumed to act in two stages, with some decisions being taken after uncertainty is resolved. In this specific case, actors must agree the rule for sharing a pie that will hold for one period (year), after which they may be able to renegotiate the agreement. Even though uncertainty over the resource to be shared in each period may decrease with time – perhaps because knowledge and ability to predict uncertain events increases – this will not affect the agreed sharing rule before the time at which agreement is to be reached. As a consequence, in as much as the uncertain policy issue is important in determining players' utility, payoffs will vary stochastically with (unpredictable) changes in this issue. Because there is no way to know what state of the world will be realised before agreement is struck, we do not expect the inclusion of uncertainty in the model to cause efficient delays, as was the case in the model developed by Merlo and Wilson (1995).

Assume that the realisation of at least one $k \in K$ is not known with certainty, and denote the uncertain policy dimension by \tilde{k} . Then let Φ denote the set of possible states of the world (SoW), with typical element φ . The stochastic element \tilde{k} and its realisations Φ are assumed to follow a specified probability distribution. Recall that, because of the constitutional features of our game, players are not able to update their beliefs about the distribution of \tilde{k} .

Refer to state $\varphi \in \Phi$ realised in period t as a state (φ, t) . For any state, let $X(\varphi, t)$ be the policy space over which players have to agree, $\mathbf{x}(\varphi, t)$ the agreed policy package, with $\mathbf{x} = (x_1, \dots, x_{\tilde{k}}, \dots, x_K)$.

In order to explore the impacts of uncertainty in one of the negotiated dimension on the bargaining strategies adopted by players, we will consider an additional constraint on players'

strategies, namely that the utility they derive from any proposal is at least equal to the utility they derive from the minimum acceptable level of the negotiated variables.

Let $\min_{i,k} = [x_{i,k}]$ be the minimum acceptable level of negotiated variable x_k for player i .

Then, a player's strategy is defined as follows:

$$[1] \quad s_i[\mathbf{x}_i, A_i, U_i(\min_{i,k})]$$

Were \mathbf{x}_i is player i 's proposal when it is his turn to make one, $A_i = \sum_{j=1}^N \omega_j U_i(\mathbf{x}_j)$ is player i 's acceptance set in the response round, and $U_i(\min_{i,k})$ is the utility player i derives from the implementation of a policy package that is minimal with respect to his requirements.

3.3 Players' equilibrium strategies

After inadmissible strategies are eliminated sequentially, players are left with a single person decision problem, and their equilibrium strategies can intuitively be characterised as in the original game by Raussler and Simon (1992): in response rounds, players will accept a proposed policy vector if and only if it yields them at least as much utility as their reservation utility for that round; in offer rounds, players will propose a vector that maximises their utility, subject to the utility of other players being at least as large as players' reservation utilities.

There is therefore a simple characterisation of the unique equilibrium strategies for a game with T bargaining rounds: when called to value another player's proposal, players only accept if the proposed policy package generates at least as much utility as their reservation utility in that bargaining round: when deciding whether or not to accept an offer, each player will compare the payoff he can get by accepting the offer with the expected payoff if he rejects it. Players' expected utility from playing in a sub-game starting at $t + \Delta$ is the ω -weighted sum of the utilities he would obtain from all other parties' proposals in that round, as every proposal that is part of the set of admissible solutions (after sequential elimination of weakly dominated strategies) is accepted (Simon et al, 2006). The reservation utility can thus be seen as the certainty equivalent of the lottery players would face if they were to reject the proposal. In the case of uncertain realisation of one of the negotiated issues, players' reservation utility will also depend on their expectations about future states of the world. In offer rounds, on the other hand, players are faced with a two-part decision problem: they maximise their utility over the set of feasible policies that provide the other players at least as much utility as their reservation utility in the following bargaining round; and they select a policy among these maximisers.

In the next section, the computational algorithm and the results of simulation exercises will be presented to illustrate the workings of the model, and assess its potential usefulness.

Solving the model

In order to explore multiple players, multiple issues negotiations within our framework, one need to specify the game in terms of players, negotiation variables, and players' utility functions. In addition, the set of feasible variables (the constraints on the negotiated variables) must be well specified. In this section, we will present the computational algorithm used to solve the model numerically and the underlying parameters. Comparative statics exercises are then carried out in order to tease out the impacts of key parameters characterising players' utility functions and, therefore, their strategies, on the equilibrium agreement.

4.1 The computational algorithm

As our negotiation game has a finite time horizon, it can be easily solved computationally. Formally, the game is solved by backward induction: first, one determines the optimal strategy of the player who makes the last move of the game. Then, the optimal action of the next-to-move player, taking the last player's action as given, is determined. This process continues until the first round of the game is reached, and all players' actions have been determined. Sub-game perfect equilibria eliminate non credible threats.

The computational algorithm developed to simulate the process of multilateral, multiple issues negotiations as described in the underlying bargaining framework is as follows: in the last period of the bargaining game – the first period of the computable model – the proponent selects a set of policy $\mathbf{x} \in X$ which maximises his utility. The other players will have to accept this proposal, because it is better for them than the disagreement policy by assumption. At round $T-1$ (round 2 of the computable model), player j will make an offer with probability ω_j , and it will be such that his utility is maximised, and the proposal yields to all other players a utility higher than their expected utility at round $T-2$ (reservation utility constraint) – which is the expected utility of moving to round $t+1$ in the bargaining game. The proposal must also be consistent with the exogenous constraints on the negotiated variables. If the solution vector to player j 's constrained maximisation also satisfies his own expected utility from proceeding to the next bargaining round, player j will propose the vector. Otherwise, he will propose a solution vector that is rejected, and the game will pass to the next round. This procedure is repeated until the proposals made by all players converge to a limit point – the equilibrium solution.

The computational model is solved recursively, by computing a series of single-person decision problems, until an acceptable degree of convergence is achieved. For the simulation, we use GAMS – General Algebraic Modelling System (McCarl, 2004)⁸.

⁸ In the GAMS code, a solution is found when $\varepsilon \leq 0.001$ – where ε is the difference between two consecutive solutions.

4.2 Players and utility functions

We define our model in the class of spatial problems, where players' utility is a declining function of the Euclidean distance between the agreed policy and players' ideal points. We adopt this perspective for the main reason that this approach is less data intensive, and better suited to represent the utility functions of designated players who do not necessarily aim at profit maximisation. Adopting the more traditional approach of equating players' preferences with their production or profit function may, in these cases, not be appropriate.

We thus represent players' preferences in the following way: the policy space $X \cup \underline{\mathbf{x}}$ consists of different locations, i.e. points in the k -dimension policy space. Each player has a most preferred location in X , called his ideal point, denoted by $\alpha = (\alpha_i)_{i \in I}$. Player i 's utility is a declining function of the Euclidean distance between a policy vector \mathbf{x} and the ideal point. Let $\mathbf{d}(\mathbf{x}, \alpha_i)$ denote the distance. In this approach, players are called idealistic or policy-seeking, and are assumed to support any policy, provided it is the best they can get. The approach often leads to compromises on all issues⁹.

For this first numerical example, let us assume that there are 5 players ($i = \{A, B, C, D, E\}$), indexed by i , who have to decide among them how to share two pies, X_1 and X_2 . For the first exercise, we will assume that both X_1 and X_2 are of known size, whereas the second exercise will introduce a random element in the size of one of the pies, X_2 . Players' utility function can be mapped as the Euclidean distance between the negotiated settlement and their most preferred location. As in the Rausser and Simon (1992) model, we will assume that the utility function of players takes the following form:

$$[2] \quad U_i(\mathbf{x}) = [\gamma_i - \mathbf{d}(\mathbf{x}, \alpha_i)]^{1-\rho_i} \quad \text{and} \quad U_i(\underline{\mathbf{x}}) = -\infty.$$

Where γ_i is a positive constant, and $\mathbf{d}(\mathbf{x}, \alpha_i)$ denotes the distance between player's ideal policy option and the proposed policy option. More specifically,

$$[3] \quad \mathbf{d}(\mathbf{x}, \alpha_i) = \sqrt{\psi + \sum_{k=1}^K \eta_{i,k} |x_{i,k} - \alpha_{i,k}|^2}$$

⁹ Alternatively, preferences could be modelled as fixed-sum game where parties are only interested in securing control over as much as possible of the available sum (money, power,...). These parties are said to be rent-seeking. Rusinowska et al. (submitted) include both rent-seeking and idealistic behaviour in one consistent model: players assign a degree of desirability and define the unacceptability set.

where $\eta_{i,n}$ denotes the strength of player i 's preferences towards issue k , ρ_i is player i 's risk aversion coefficient – with $\rho_i \in 0,1$. Finally, ψ is a positive constant to avoid indeterminacies of the square root; this constant is the same for all players, and has the only effect of scaling the utility function – thus the relative differences are not affected.

In line with the description of the model, players will accept a proposal only if it yields them at least as much utility as their expected continuation payoff. For any proposal \mathbf{x} , players' participation constraint takes the following form:

$$[4] \quad U_i(\mathbf{x}_T) \geq U_i(\underline{\mathbf{x}}) = 0 \text{ and}$$

For $t = T$, and for $t \leq T$,

$$[5] \quad U_i(\mathbf{x}_t) \geq EU_i = \sum_{j=1}^N \omega_j U_i(\mathbf{x}_{t+1}^j)$$

That is, the utility player i would derive from accepting the proposal by player j at time t ($U_i(\mathbf{x}_t^j)$) must be at least as large as player's expected utility from rejecting the proposal and moving to time $t+1$. As stated in previous sections, player i 's expected utility (or acceptance set) is equal to the sum of all the possible payoffs he could get in the next stage of the game, determined by the proposal of all the remaining players including himself, weighted by the proposer's access probability, ω_i , with $\sum_i \omega_i = 1$. Table 1 reports the values that the utility parameters take for the numerical simulations.

The allocation of the pie is constrained by the total quantity available to share, thus the sum of the quantities allocated to each player cannot exceed the total quantity available:

$$[6] \quad \sum_N x_{i,1} \leq \overline{X}_1 \text{ and } \sum_N x_{i,2} \leq \overline{X}_2$$

Where $x_{i,1}$ and $x_{i,2}$ are the amount of X_1 and X_2 allocated to player i respectively, \overline{X}_1 and \overline{X}_2 are the respective sizes of the pies, determined exogenously.

The vector of access probability – establishing the order in which players are called upon to make a proposal – is determined exogenously. As observed in the previous chapters, there are several ways in which the access probability of players can be determined. In our simulations, we will initially assign to players the same access probabilities.

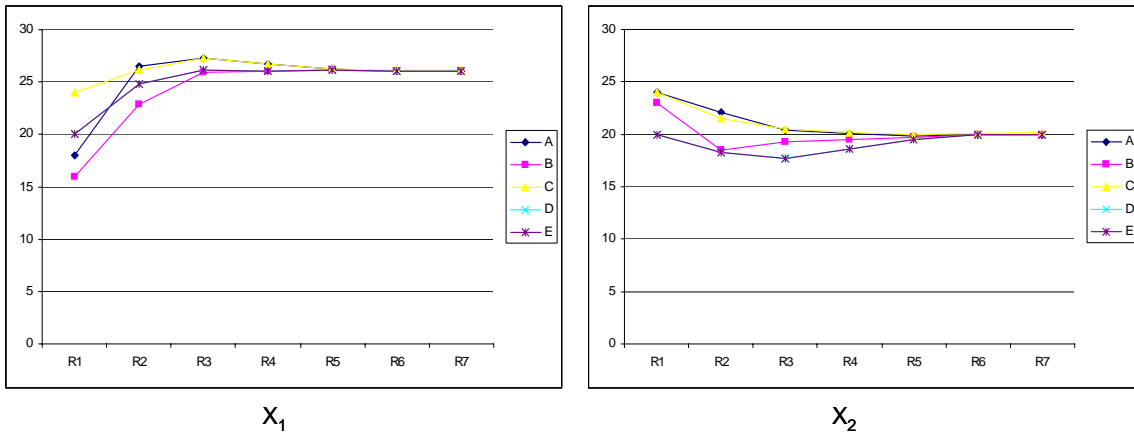
Table 1: Parameters of the utility function in the baseline equilibrium

Player (i)	ω_i	Ψ	γ_i	$\alpha_{n,i}$		$\eta_{k,i}$		ρ_i
				$k = x_1$	$k = x_2$	$k = x_1$	$k = x_2$	
A	0.2	100	100	18	24	0.3	0.7	0.5
B	0.2	100	100	16	23	0.4	0.6	0.5
C	0.2	100	100	24	24	0.6	0.4	0.5
D	0.2	100	100	20	20	0.5	0.5	0.5
E	0.2	100	100	20	20	0.5	0.5	0.5

4.3 The baseline equilibrium

Let us present the main outputs of our numerical solution of the bargaining model. Let us first focus on the baseline equilibrium, i.e. the equilibrium computed using the parameters of Table 1. In the next section, we will perform a careful sensitivity analysis to identify the factors that most influence the equilibrium of the dynamic bargaining game. In our baseline, players' proposals converge by the 7th iteration, as shown in Figure 1. The rate of decline of the proposals is smaller for the last two players, who are also able to extract a higher surplus at the equilibrium.

Figure 1: Convergence of players' proposals – x_1



As predicted by the characterisation of players' strategies, if the final round of the negotiation is reached (the first round of our computational model), each player, when selected to be the proposer, will propose an allocation such that his utility is maximised – that is, he will propose his ideal point. The respondents will have to accept it, as the proposal will yield them a strictly higher utility than the disagreement policy. This is shown in Table 2: when A is the proposer, the values of x_1 and x_2 correspond to player A's ideal point – and the same applies to all the other players. As a consequence,

the utility derived by each player in the final round T , when they are selected to be the proposer, is at its highest.

When responding to an offer, players will compare the utility they would attain by accepting the proposal with what they expect to get by rejecting it. This is the ω -weighted sum of the expected utility when each of the other player is a proposer, denoted by EU in Table 2. In round $T - 1$, the proposer will maximise his utility, subject to the physical constraints on the cake being satisfied, but with an additional constraint – that the proposal yields his opponent a utility at least equal to their expected continuation payoffs. As shown in Table 2, therefore, when asked to be proposers in the second to last round of the game (the second round of our computational model), respondents will ask for a lower allocation of the cake, compatible with the participation constraint of others. The game continues in this fashion until a satisfactory degree of convergence is ensured.

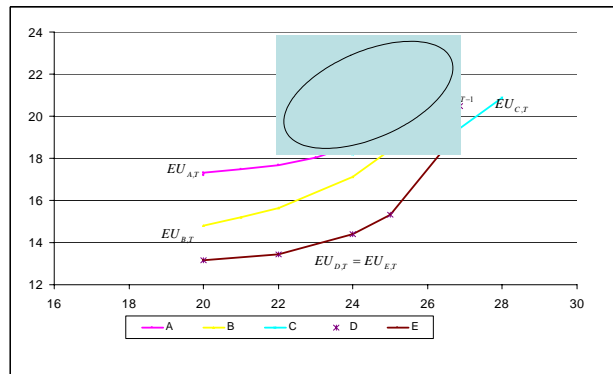
Table 2: Baseline equilibrium

Simulation round (negotiation round)	Player	A	B	C	D	E
	Proposer					
Round 1 (T)	A	9.486833	9.331099	9.437363	9.406866	9.406866
	B	9.370103	9.486833	9.43356	9.397909	9.397909
	C	9.391632	9.359616	9.486833	9.432477	9.432477
	D	9.393666	9.349774	9.4517	9.486833	9.417047
	E	9.393666	9.349774	9.4517	9.417047	9.486833
EU		9.40718	9.375419	9.452231	9.428226	9.428226
Round 2 (T-1)	A	9.426226	9.375419	9.460392	9.442688	9.442688
	B	9.40718	9.410532	9.46129	9.429141	9.429141
	C	9.40718	9.375419	9.472804	9.441784	9.441784
	D	9.413648	9.375419	9.465461	9.453845	9.436373
	E	9.413648	9.375419	9.465461	9.436373	9.453845
EU		9.413576	9.382442	9.465082	9.440766	9.440766
Round 3 (T-2)	A	9.401423	9.382442	9.465082	9.44573	9.44573
	B	9.413576	9.373819	9.465277	9.440766	9.440766
	C	9.413576	9.382442	9.457437	9.445996	9.445996
	D	9.418995	9.382442	9.468582	9.432993	9.441387
	E	9.418995	9.382442	9.468582	9.441387	9.432993
EU		9.413313	9.380717	9.464992	9.441375	9.441375
Round 4 (T-3)	A	9.4044	9.380717	9.464992	9.443828	9.443828
	B	9.413313	9.37314	9.465079	9.441375	9.441375
	C	9.413313	9.380717	9.459902	9.44395	9.44395
	D	9.416361	9.380717	9.466964	9.437712	9.441375
	E	9.416361	9.380717	9.466964	9.441375	9.437712
EU		9.41275	9.379202	9.46478	9.441648	9.441648
Round 5 (T-4)	A	9.406893	9.379202	9.46478	9.44225	9.44225
	B	9.412859	9.373683	9.46478	9.441648	9.441648
	C	9.41275	9.379202	9.461691	9.442393	9.442393
	D	9.413646	9.379202	9.465268	9.439272	9.441648
	E	9.413646	9.379202	9.465268	9.441648	9.439272
EU		9.411959	9.378098	9.464358	9.441442	9.441442

Round 6 (T-5)	A	9.410017	9.378131	9.464358	9.441442	9.441442
	B	9.412158	9.376108	9.464358	9.441442	9.441442
	C	9.411959	9.378409	9.463639	9.441979	9.441979
	D	9.412158	9.378131	9.464358	9.440472	9.441442
	E	9.412158	9.378131	9.464358	9.441442	9.440472
EU		9.41169	9.377782	9.464214	9.441355	9.44135
Round 7 (T-6)	A	9.410988	9.37792	9.464214	9.441355	9.441355
	B	9.411919	9.376963	9.464214	9.441355	9.441355
	C	9.41169	9.378152	9.46428	9.441853	9.441853
	D	9.411919	9.37792	9.464214	9.440831	9.441355
	E	9.411919	9.37792	9.464214	9.441355	9.440831

It is interesting to note that players, when selected to be proposers, face different binding constraints. Consider, for instance, the case in which player C is selected to be the proposer in round $T - 1$. Figure 2 depicts the situation that player C is facing: his proposals to the other players are indicated by the $x_{i,T-1}^C$'s, whereas the solid lines denoted by $EU_{i,T}$ are players' participation constraints in that round – that is, combinations of x_1 and x_2 that yield to player i the same payoff as his expected payoff from the lottery in round T , generated by players' proposals in the final round. In other words, it is the expected payoff to each player, should they reject the proposals in round $T - 1$, and should the final round T be reached.

Figure 2: Players' participation constraints in round $T-1$ when player C is the proposer



When player C is determining the proposals to make in the final round of the game, he will be faced with two binding constraints – namely, the expected utilities of players A and B – and the zone of agreement is represented by the shaded area. It is clear that the participation constraints of the remaining two players are slack – thus they benefit from a higher utility in round $T - 1$ when C is the proposer, thanks to the bargaining strength of players A and B which force player C to make a proposal less favourable to himself.

This can partly be explained by the fact that the ideal point influences players' utilities: those players with lower ideal points will gain more than the others. The underlying reason is simple (see Simon et al, 2003): when players have very different preferences from the others, their participation constraints become binding on the other players.

Sensitivity Analysis


Despite being simple, the above exercise illustrates how the deadline effect propagates through the model – that is, how changes in the last round of the bargaining game build up to magnify as we move backwards in the game. Let us now focus on other important features of the bargaining game and on their role in shaping the equilibrium outcome. First, let us analyse the role of players' relative bargaining power. To analyse it, we perform a second round of simulations which are identical to the baseline case, with the exception of players' bargaining power. In particular, we let us assume that bargaining power shifts from player *A* to player *B*, keeping the power of other players constant, as shown in Table 3.

Table 3: Sensitivity analysis – changing bargaining power

Player (<i>i</i>)	ω_i	
	Initial location	Perturbed location
A	0.2	0.1
B	0.2	0.3
C	0.2	0.2
D	0.2	0.2
E	0.2	0.2


Intuitively, one would expect that, while the player with a higher access probability increases his equilibrium utility level, the other players experience a reduction in utility. The effect of increasing a player's bargaining power is to tighten the constraint the other players face when proposing an allocation, as shown in Figure 3, where the participation constraint for player *B* is higher in the case where bargaining power is shifted than in the baseline case. Similarly, the effects of increasing a players' bargaining power is to relax the participation constraints of player *A* in the penultimate round of negotiation, while leaving the other players' participation constraints substantially unchanged. Thus, if the negotiations proceed to round $T - 1$, player *B* may be able to extract a larger surplus than in the baseline case.

Figure 3: Expected utility constraints in round T-1- the effect of shifting bargaining power.



Consider Figure 4. Each cluster of bars decomposes the change in players’ expected utilities, as access is shifted from player A to player D – keeping other players’ bargaining power constant. In round $T - 1$, the utility that player A gains from his own proposal in the perturbed situation – compared to the baseline utility – is lower: this is because the participation constraint of player B on player A is binding in the penultimate round in the baseline case (see Table 2): increasing the bargaining power of player B will necessarily lead to a tighter expected utility constraints, thus forcing player A to make a proposal which is less favourable to himself than would otherwise be the case. At the same time, player B gains substantially, independently of who is proposing in the penultimate round of the negotiation process.


Figure 4: Change in utilities from shifting bargaining power – round T-1



At the equilibrium, the effects of shifting bargaining power from one player to another, keeping the other constant, may not be linear. The implication is that the ultimate effect of shifting bargaining power cannot be predicted, as this is not linearly related to the final outcome of the negotiation, but critically depends on other constitutional factors – such as decision rules, players’ preference parameters, the relative distance of their ideal points, so on and so forth. This is shown by

the results of our simulation: the final impact depends on the bargaining design. Figure 5 shows the evolution of players' participation constraints as we move backward in the game – that is, the difference between players' expected utilities between the baseline case and the perturbed case.

Figure 5: Changes in players' participation constraints



It is clear that, through various interacting forces, the disadvantage of player *A* decreases as we move along the game, while the advantage of player *B* decreases. The cause of this effect is to be found in the preference structure and utilities of the players: the distance between player *A* and player *B*'s ideal points is relatively small, as compared to the other two players (see Table 4). Thus, in the longer run, player *A* is able to benefit from the increased bargaining power of player *B*.

Table 4: Distance between players' ideal point

	A	B	C	D	E
A	0	2.24	6.00	4.47	4.47
B	2.24	0	8.06	5.00	5.00
C	6.00	8.06	0	5.66	5.66
D	4.47	5.00	5.66	0	0
E	4.47	5.00	5.66	0.00	0.00


Let us now look at the impact of restricting the values that the negotiated variables can take. If the predictions of theoretical models are correct, we would expect to see that players have a lower utility from the negotiations, as the bargaining space – and as a consequence the gains from trade – is restricted. The new upper bounds for the negotiated variables are presented in Table 5.

Table 5: Perturbed constraints on the negotiated variables

		$n = 1$	$n = 2$
\bar{X}_n	Baseline	130	100
	Perturbed	130	80


Figure 6 shows the changes in utility level for each player when they are selected to be the proposers in the respective negotiation rounds – that is, it is the difference in utilities they enjoy when proposing an agreement in the restricted and baseline cases. It is clear that all players experience a decrease in utility as compared to the baseline case. The decrease is however not uniformly affecting all players: those who have stronger preferences for the restricted policy issue (that is, those with a higher ideal point for X_2 , players A , B , and C) suffer more from this restriction than the other players. The decline in utility is mitigated by the different weights that individual players assign to X_2 relative to X_1 , as indicated by the different values of $\eta_{i,k}$ in Table 1: thus, players A and B , who have a stronger preference towards X_2 as compared to X_1 , suffer a loss higher than player C , who has a higher preferred point for X_2 , but assigns a low weights to this variable relative to the previous two players. This result of the simulation exercise is in line with both the theoretical findings of non-cooperative bargaining theory and the applications of non-cooperative bargaining models to water negotiations (see Carraro et al., 2005, and Carraro et al., 2007). Furthermore, more iterations are needed before a limit point equilibrium solution is found, indicating the increased difficulties in finding a compromise allocation.

Figure 6: Changes in players' utilities from restricting the issue space



Interestingly, the players with a higher ideal point for the restricted issue will “bargain harder” in the last rounds of the negotiation game, and require a higher share of the total resource available for themselves. This effect is shown in Figure 7: should the final round of the negotiation game be reached, the first three players, when selected to be proposers, will ask for themselves a higher share of X_2 . The effect decreases for player C as the game proceeds backward, because of the larger opportunity that this player has to compensate for losses in X_2 with higher X_1 . These results are robust to further restriction in the issue space.

Figure 7: Change in players' proposed shares – x_2



Consider now the case in which one of the player's preferences towards the two negotiated variables changes in such a way that he now strongly prefers satisfying his ideal quantity of one negotiated variable relative to the other. Intuitively, one would expect the equilibrium shares of this player to change so that his equilibrium quantity of the strongly preferred variable is higher than in the baseline simulation exercise.


In the simulation exercise, we change the relative weights that player D assigns to X_2 relative to X_1 , as shown in Table 6.

Table 6: Changing the relative importance of the negotiated variables.

Player D	$\eta_{D,1}$	$\eta_{D,2}$
Initial	0.5	0.5
Perturbed	0.2	0.8


As expected, player D will require a higher share of X_2 relative to the baseline case (see Figure 8).

Figure 8: Variation in player D proposals for x_2



What is of interest is that his stronger position affords player *D* a stronger bargaining position, and his participation constraints in round $T - 1$ tightens, while the participation constraints of the remaining players are substantially unchanged (Figure 9, $T - 1$). This effect is preserved through the backward induction process, and in equilibrium player *D* will enjoy a higher utility level (Figure 9, Equilibrium). In this model, there seem to be two sources of bargaining power: first of all, players' access – which, however, is neither linearly nor monotonically related to players' equilibrium payoffs – and players ideal points – both in terms of their magnitude and relative importance.

Figure 9: Changes in players' participation constraints in round $T-1$



Finally, it is interesting to note that, in the long run, a relatively large change in players' weighting of the negotiated variables leads to utility levels that are significantly higher for that player, but leave the expected utilities of other players substantially unchanged.

The role of uncertainty

One of the key aspects of negotiation processes is uncertainty over the size of the negotiated variables (the size of the pie). Let us analyse how this type of uncertainty affects the agreement and players' utilities. In the numerical analysis, we look at the impact of introducing a random component in the constraint function for \bar{X}_2 . The new constraint for this variable will thus take the following form:


$$[7] \quad \sum_N x_{i,2} \leq \tilde{X}_2$$

where \tilde{X}_2 is an uncertain component of the size of the pie to be divided. In the case of negotiations on water availability, for instance, the total quantity available depends in part on precipitation levels, which, however, cannot be predicted with certainty.

To demonstrate the relationship between the introduction of uncertainty in the realisation of one of the negotiated variables and the frequency of different solution, we report the results of a Monte Carlo experiment, in which we solve the model for 100 randomly drawn values of \tilde{X}_2 , assuming an exogenously specified underlying probability distribution for the unknown term. Using random inputs, the deterministic model is essentially turned into a random model.

The choice of the underlying distribution to simulate random sampling may impact the results of the simulations. For the numerical example, we will assume that the random component of \tilde{X}_2 , is drawn from a gamma probability distribution¹⁰, with shape parameter 13 and scale parameter 8.5. The corresponding mean and standard deviations are, respectively, 104.4 and 29.3. Figure 10 shows the frequency distribution of the realised values of \tilde{X}_2 , together with some basics statistics.

Figure 10: Distribution of X_2



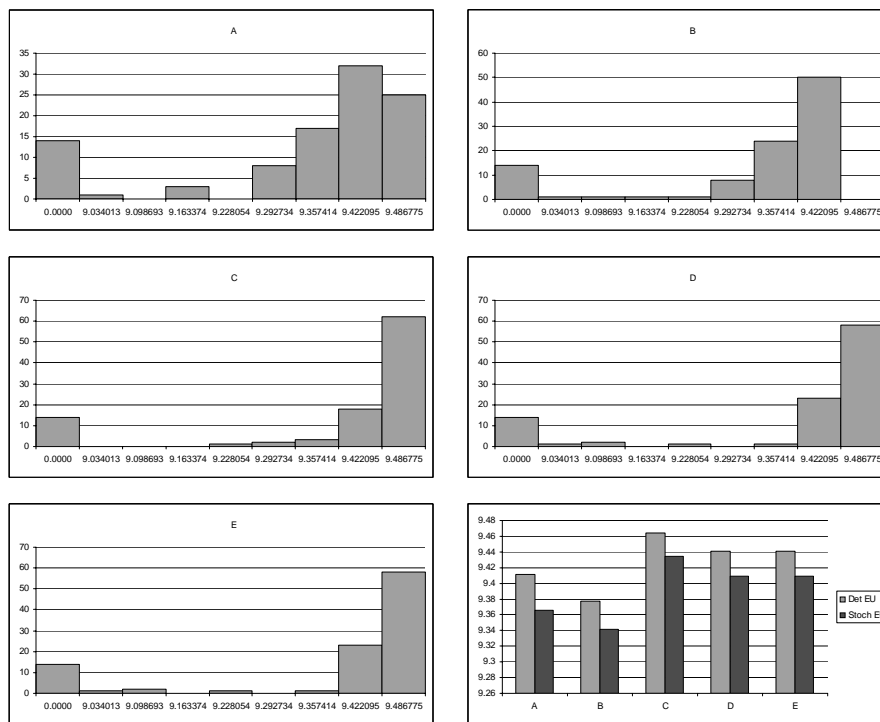
For each of the 100 sampled bargaining problems, we examine the emerging strategies of players, as well as the equilibrium solutions and utilities, for 100 bargaining rounds.

¹⁰ There are various types of probability distribution one could use: the normal distribution is applicable to variables whose values are determined by an infinite number of independent random events; very rare events are best represented using the Poisson distribution. The normal distribution is symmetrical around the mean and, in general, it is used when (i) there is a strong tendency for the variable to take a central value; (ii) positive and negative deviations from the central value are equally likely; and (iii) the frequency of deviations falls off rapidly as the deviations become larger. The gamma distribution is, on the other hand, widely used in engineering to model continuous variables that are always positive and have a skewed distribution.

From the results, it would appear that *it takes longer to find an agreement when uncertainty over the resource is included*. Moreover, in 14% of the cases, a feasible agreement cannot be achieved – that is, the equilibrium offers of players are not compatible with the resource constraint.

The first 5 histograms of Figure 11 show the frequency distribution of players’ utilities when the size of negotiated variable is uncertain – so, for instance, player *D* and *F* experience more frequently low utility levels when all the players entertain the possibility of variations in the size of one pie. The variability of players’ equilibrium utilities differs among the players, while for all of them the equilibrium (limit) utility is lower in the stochastic case than in the deterministic case, as shown by the lower right quadrant of Figure 11.

Figure 11: Frequency distribution of players’ equilibrium utilities



Of course, the deterministic case assumes a known realisation of the pie. But what happens to players’ utilities if this belief is mistaken? Would they be better off by considering this possibility in their bargaining strategy?

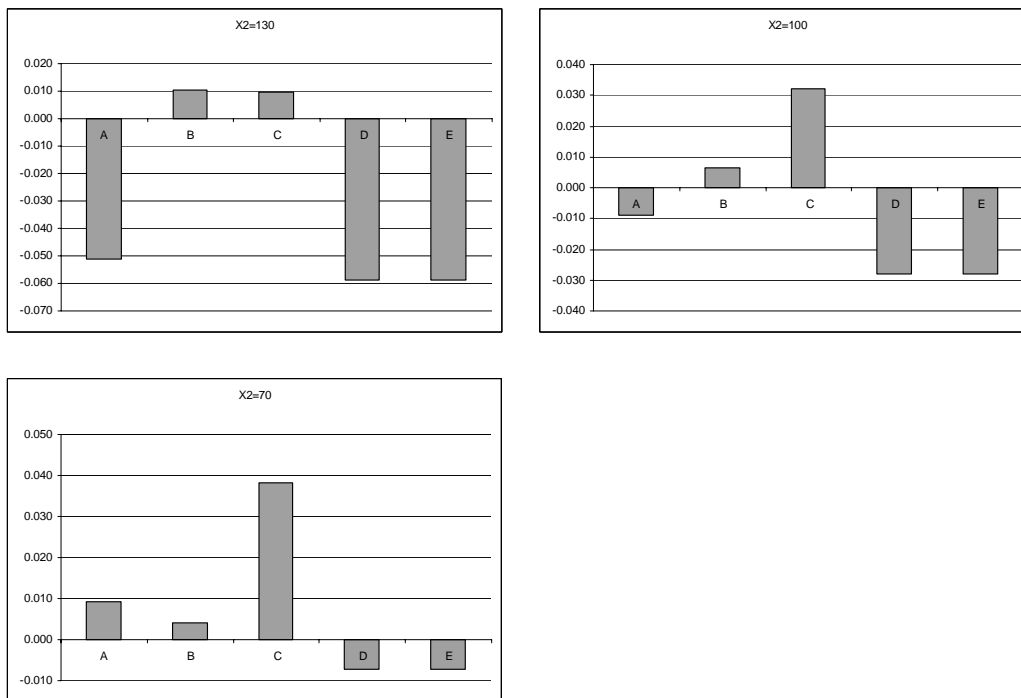
In order to assess the ex post efficiency of the different negotiation frameworks, let us compare the utilities that player would derive, in equilibrium, from the negotiated agreement under the deterministic and stochastic case. Table 7 below reports the equilibrium shares agreed upon in the deterministic and stochastic models.

Table 7: Equilibrium shares

X1	A	B	C	D	E
Deterministic	0.2002	0.2002	0.2013	0.2002	0.2002
Random	0.2026	0.1986	0.2027	0.2018	0.2018
X2					
Deterministic	0.2004	0.2001	0.2014	0.1985	0.1985
Random	0.1997	0.1997	0.2015	0.1997	0.1997

As shown in the first left hand side quadrant of Figure 12, players do, in general, better in the deterministic type of model as compared to the stochastic type. However, as the resource available begins to shrink, players' utility increases, in general, when they take into account the uncertainty surrounding the realisation of the surplus. The explanation is intuitive: *as players begin to account for uncertainty in their strategy, they will try to negotiate harder, expecting a higher share of the surplus, in order to increase their chance of coming closer to their ideal*¹¹.

Figure 12: Changes in players' equilibrium utility – deterministic vs. stochastic



¹¹ Note that, because of the construction of our preference function, an excess allocation of water is a punishment for players. This may not be realistic in some circumstances, as discussed more in details in the concluding section.

Conclusions

The model proposed in this paper attempts to simulate the process of negotiation among multiple players, who have to decide on how to share a surplus of fixed size. In this context, negotiation rules are simulated through an offer and counteroffer procedure. Players have payoff functions that depend on the share of the surplus that they can secure for themselves – with different negotiated variables having different importance for each player, thus generating space for tradeoffs among them. Furthermore, players have varying access probabilities, which signal the relative strength at the bargaining table and thus influence the equilibrium agreement.

Through a series of simulations in which five players negotiate over the respective shares of two cakes, we have examined the emerging equilibrium agreements and their characteristics. What can be inferred about these problems by applying numerical simulation?

First of all, the results conform to expectations when there is no uncertainty over the negotiated variables. As in the Rausser-Simon model and its applications, increasing the access probability of a player will yield outcomes that are more favourable for the “more powerful” player, but also to players with similar preferred positions. Convergence of the solution is attained in few iterations of the model – which can in part address some of the critiques moved to backward induction, as there is scepticism of long and involved inductive chains.

This result does no longer hold when we restrict significantly the range of admissible values for the negotiated variables. In fact, restricting the size that the negotiated variables can take reduces the opportunities for trade, yielding potentially lower utilities to all players. In some cases, excessively reducing the boundaries of the negotiated variable may shrink the bargaining space so much that no zone of agreement remains. Should this result emerge when exploring a real problem using this framework, it would be advisable to attempt changing the decision rule – from unanimity to qualified majority, for instance.

Importantly, the effect of bargaining power on the equilibrium agreement is non linear, but rather evolve in complex way through the process of backward induction. The effect of shifting access depends crucially on other constitutional factors with which it interacts – such as decision rules, players’ preference parameters, the relative distance of their ideal points, so on and so forth. Thus, there are synergies among players or issues that affect the ultimate impact of bargaining, contrary to the assumption of the standard Nash games.

Finally, uncertainty over one of the negotiated variables crucially affects the equilibrium outcome and the players’ strategies. Our main results are:

- (i) when uncertainty is introduced, the negotiation takes, on average, longer (14 rounds as opposed to 7 rounds in the deterministic case);

- (ii) in some cases, players' strategies do not even converge to a feasible solution – that is, players' offers crystallise on values that are not compatible with the resource constraint for neither variables;
- (iii) explicitly accounting for uncertainty in the realisation of the surplus leads, under some circumstances, players to bargain harder: ex post, they are better off only when the realisation of the surplus is low, as compared to the deterministic case.

These results are in line with intuition and with previous results of similar models: they therefore lend support to the hypothesis that non-cooperative bargaining is a useful framework for exploring negotiation processes and players' strategic behaviour. Applying non-cooperative bargaining theory can provide some useful insights, based on formal models, as to which factors influence to a significant extent players' strategies and, as a consequence, the resulting equilibrium agreement policy.

The value added of exploring management problems within a non-cooperative bargaining framework lies in the ability of the approach to help finding politically and socially acceptable compromise. The proposed model can find several applications, in particular, in the field of natural resource management – where conflicts over how to share a resource of a finite size are increasing.

Acknowledgements

The authors are grateful to Ariel Dinar, Fioravante Patrone and above all Carmen Marchiori for helpful suggestions and discussions. Comments and remarks from Carlo Giupponi, Ignazio Musu and Paolo Rosato are also gratefully acknowledged. All remaining errors are obviously ours.

References

- Adams, G., Rausser, G., & Simon, L. (1996). Modelling multilateral negotiations: an application to California Water Policy, *Journal of Economic Behaviour and Organization*, 30, 97-111.
- Alesina, A., Angeloni, I., & Etro, F. (2001). *The political economy of unions*, NBER Working Paper.
- Bac, M., & Raff, H. (1996). Issue-by-issue negotiations: the role of information and time preferences, *Games and Economic Behavior*, 13, 125-134.
- Busch, L.A., & Horstmann, I. (1997). Signaling via agenda in multi-issue bargaining with incomplete information, mimeo.
- Carraro, C., Marchiori, C., & Sgobbi, A. (2005). *Applications of Negotiation Theory to Water Issues*, World Bank Policy Research Working Series, 3641, The World Bank.
- Carraro, C., Marchiori, C., & Sgobbi, A. (2007). Negotiating on Water. Insights from Non-cooperative Bargaining Theory, *Environment and Development Economics*, 12(2). pp 329-349.

- Carraro, C., & Siniscalco, D. (1997). *R&D cooperation and the stability of international environmental agreements*, in *International Environmental Negotiations: Strategic Policy Issues* edited by Carraro, C.
- Cesar, H., & De Zeeuw, A. (1996). Issue linkage in global environmental problems, in *Economic Policy for the Environment and Natural Resources*, edited by Xepapadeas, A.
- Chae, S., & Yang, A. (1988). The unique perfect equilibrium of an N-person bargaining game, *Economic Letters*, 28, 221-223.
- Chae, S., & Yang, A. (1994). A N-person pure bargaining game, *Journal of Economic Theory*, 62, 86-102.
- Diermeier, D., Eraslan, H., & Merlo, A. (2004). *Bicameralism and Government Formation*, CTN 9th Workshop 04.
- Eraslan, H., & Merlo, A. (2002). Majority rule in a stochastic model of bargaining, *Journal of Economic Theory*, 103, 31-48.
- Fatima, S., Wooldrige, M., & Jennings, N. R. (2003). Optimal agendas for multi-issue negotiation, mimeo, University of Southampton.
- Fershtman, C. (1990). The importance of agenda in bargaining, *Games and Economic Behavior*, 2, 224-238.
- Folmer, H., van Mouche, P., & Ragland, S. E. (1993). Interconnected games and international environmental problems, *Environmental Resource Economics*, 3, 313-335.
- Gneezy, U., Haruvy, E., & Roth, A. E. (2003). Bargaining under a deadline: evidence from the reverse ultimatum game, *Games and Economic Behaviour*, 45, 347-368.
- Herrero, M. (1985). *A strategic theory of market institutions*, London: London School of Economics.
- In, Y. & Serrano, R. (2003). Agenda restrictions in multi-issue bargaining (II): unrestricted agendas, *Economics Letters*, 79, 325-331.
- Inderst, R. (2000). Multi-issue bargaining with endogenous agenda, *Games and economic behavior*, 30, 64-82.
- Jun, B. H. (1987). *A structural consideration on 3-person bargaining*, Department of Economics, University of Pennsylvania.
- Krishna, V., & Serrano, R. (1996). A model of multilateral bargaining, *Review of Economic Studies*, 63, 61-80.
- Kultti, K., & Vartiainen, H. (2004). *A non-cooperative solution to the bargaining problem*. Working paper.
- Ma, C., & Manove, M. (1993). Bargaining with deadlines and imperfect player control, *Econometrica*, 61, 1313-1339.
- Manzini, P., & Mariotti, M. (2003). A bargaining model of voluntary environmental agreements, *Journal of Public Economics*, 87(12), 2725-2736.
- Manzini, P., & Mariotti, M. (2005). Alliances and negotiations, *Journal of Economic Theory*, 121(1), 128-141.
- McCarl, B. A. (2004). *GAMS User Guide: 2004. Version 21.3*, GAMS Development Corporation.
- Merlo, A. (1997). Bargaining over government formation in a stochastic environment, *Journal of Political Economy*, 105, 101-131.
- Merlo, A., & Wilson, C. (1995). A stochastic model of sequential bargaining with complete information, *Econometrica*, 63, 371-399.

- Myerson, R. B. (1978). Refinements of the Nash equilibrium concept, *International Journal of Game Theory*, 15, 133–154.
- Rausser, G., & Simon, L. (1992). *A non cooperative model of collective decision making: a multilateral bargaining approach*, Department of Agricultural and Resource Economics, University of California, Berkeley.
- Rubinstein, A. (1982). Perfect Equilibrium in a Bargaining Model, *Econometrica*, 50, 97-110.
- Rusinowska, A., de Swart, H., & van der Rijt, J. W. A new model of coalition formation, *Submitted for publication*.
- Shaked, A. (1986). A three-person unanimity game, mimeo.
- Simon, L., Goodhue, R., Rausser, G., Thoyer, S., Morardet, S., & Rio, P. (2003). *Structure and power in multilateral negotiations: an application to French Water Policy*, edited.
- Simon, L., Goodhue, R., Rausser, G., Thoyer, S., Morardet, S., & Rio, P. (2006). *Structure and power in multilateral negotiations: an application to French water policy*, paper presented at 6th Meeting of Game Theory and Practice, Zaragoza, Spain, 10-12 July 2006.
- Ståhl, I. (1972). *Bargaining theory*, Stockholm School of Economics.
- Thoyer, S., Morardet, S., Rio, P., Simon, L., Goodhue, R., & Rausser, G. (2001). A bargaining model to simulate negotiations between water users, *Journal of Artificial Societies and Social Simulation*, 4.