

ECONOMIC RESEARCH REPORTS

***A STOCHASTIC MODEL OF
SEQUENTIAL BARGAINING WITH COMPLETE
INFORMATION AND TRANSFERABLE UTILITY***

BY

***Antonio Merlo
and
Charles Wilson***

RR # 92-34

October, 1992

**C. V. STARR CENTER
FOR APPLIED ECONOMICS**



**NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, N.Y. 10003**

A Stochastic Model of Sequential Bargaining
with Complete Information and Transferable Utility

by

Antonio Merlo
University of Minnesota

and

Charles Wilson
New York University

October 20, 1992

The authors gratefully acknowledge the support of the C.V. Starr Center for Applied Economics. We also wish to thank Roberto Chang, Prajit Dutta, and the participants of the microeconomics workshop at New York University for helpful comments.

ABSTRACT

We consider an n-person sequential bargaining model in which both the surplus to be allocated as well as the identity of the proposer in each period follows a stochastic process. Assuming only one agent makes an offer in each period and agreement must be unanimous, we provide characterizations of both the sets of subgame perfect payoffs and stationary subgame perfect payoffs for the case where utility is transferable among the players. Our analysis incorporates several models of sequential bargaining which build on the work of Rubinstein (1982). We also investigate conditions under which agreement will be delayed and extend existing results for the uniqueness of subgame perfect and stationary subgame perfect payoffs to the stochastic game. JEL Classification Numbers: C73, C78. Key Words: Non Cooperative Bargaining, Dynamic Games, Stochastic Games.

1. Introduction

Existing models of sequential bargaining, building on the work of Rubinstein (1982) and Stahl (1972), typically have the feature that the discounted value of the surplus to be divided decreases in a deterministic fashion with each rejected offer. Although this assumption may be appropriate for a large class of bargaining problems such as the classic problem of "dividing the dollar", it may not be adequate for many other bargaining contexts. In those situations in which the agents are bargaining over complex agreements and a significant period of time is required to formulate counter offers, as for example in negotiations over labor contracts or political agreements, an efficient allocation of the surplus may be very sensitive to aspects of the environment which may change during the negotiating period. Consequently, given practical limitations on the kind of state contingent agreements which may be reached, the "cake" of the standard bargaining models as well as the bargaining power of agents may evolve over time according to a nontrivial stochastic process. In such an environment, the players may find it in their interests to delay agreement until key elements of uncertainty are resolved.

In this paper we study an n -player sequential bargaining game with complete information and transferable utility in which both the surplus to be allocated as well as the identity of the proposer follows a stochastic process. In each period a state is realized which determines the total utility to be allocated in that period if an agreement is reached as well as the order in which the players move. The first player to move in any period may either propose an allocation or pass. If he proposes an allocation, each of the remaining players in turn accept or reject the proposal. If any player rejects the proposal, a new state is realized and the process is repeated until some proposed allocation is unanimously accepted.

We provide a general characterization of both the subgame perfect and stationary subgame perfect payoffs in terms of the fixed points of operators on the set of payoff functions. Our approach is essentially an extension of the method adopted by Shaked and Sutton (1984a) and Binmore (1987) in their analyses of nonstochastic bargaining games. However, because both the equilibrium allocation and the length of bargaining are random variables, we are able to address a number of issues which cannot be analyzed in models where the environment follows a deterministic path. For instance, we may investigate the conditions under which agreement is delayed even when the equilibrium is unique. How the processes determining the proposer and the cake size affect the bargaining power of the individual players may also be investigated. Not only may we change the order or the frequency in which the players move, but we may also explore how the equilibrium payoffs are affected when the identity of the proposer is correlated with the size of the cake.

In this paper, we restrict our analysis to the case of transferable utility. That is, we assume that the cake has a symmetric linear boundary and that agents have a common discount factor so that the total payoff is always well defined. The complementary analysis for the more general model with nontransferable utility is provided in Merlo and Wilson (1992). There we establish very similar characterizations of the sets of subgame perfect and stationary subgame perfect payoffs, although some of the more specific results such as the uniqueness of stationary subgame perfect equilibria must be qualified.

Since the game includes the standard model of alternating offers, our analysis synthesizes and generalizes much of the work on sequential bargaining with complete information. For instance, the game incorporates Binmore's

(1987) model with a random selection of proposers. Also, because we impose no restrictions on the process determining cake size, the case of exogenous risk of breakdown of negotiations (see e.g. Binmore, Rubinstein and Wolinsky (1986)) is also included. Although not explicitly incorporated into our framework, we could also modify the game so that players may take up outside offers after rejecting a proposal as in Shaked and Sutton (1984b) and Binmore, Shaked and Sutton (1989) without changing the nature of our basic results. The n-player generalization of the Rubinstein model studied by Herrero (1985) and Moulin (1986) is also included in our framework.

There are, however, limits to the generality of the model imposed by our restrictions that only one agent makes an offer in each period and agreement be unanimous. This rules out games in which simultaneous offers are allowed as in Chatterjee and Samuelson (1990) and Stahl (1990), as well as continuous time bargaining models with endogenous order and timing of moves as in Perry and Reny (1989) and Sakovics (1990).¹ We also leave unexplored the consequences of assuming different voting rules other than unanimity (see Baron and Ferejohn (1989)). Nor can we easily accommodate the extensive forms studied by Chae and Yang (1988) and Krishna and Serrano (1991), in which the multilateral bargaining procedure is reduced to a series of bilateral negotiations and acceptance leads to the exit of the accepting players from the game. In contrast to the game we study, the SP equilibrium in these games is unique for any number of players.

¹ We could allow for simultaneous responses, but that introduces additional equilibria in games with more than two players since more than one player may simultaneously reject an acceptable offer. Haller (1986) explores some of the implications of this framework.

Finally, our analysis focuses exclusively on the case of complete information, in contrast to much of the recent literature which focuses on the possibility of delaying agreement as a device to reveal information about one's opponent. (See, e.g., Admati and Perry (1987) and Sobel and Takahashi (1983). Osborne and Rubinstein (1990) provide a more extensive survey of this literature.) We also investigate the possibility of delay in this paper, but the underlying motivation is quite different. In our framework, players are not trying to discover the true preferences of their opponents, but rather are seeking more efficient agreements.

The paper proceeds as follows. In Section 2, we define the game and the fundamental equilibrium concepts to be employed, followed in Section 3 by a detailed example to illustrate some of the key points in the analysis to follow. Section 4 contains the heart of the analysis. We begin by establishing a basic characterization of the set of subgame perfect payoffs supported by stationary strategies (SSP payoffs). From this characterization, we demonstrate the uniqueness and efficiency of any SSP payoff. We characterize the conditions under which agreement is delayed and provide sufficient conditions for agreement to follow a "reservation rule" with respect to the realized cake size.

In Section 6, we drop the stationarity restriction and turn our attention to the set of all subgame perfect payoffs. Although the possibility of nonstationary threats complicates the analysis, we are still able to provide a characterization of the subgame perfect (SP) payoffs in terms of the fixed points of an operator on the set of payoff functions. Using this characterization, we show that, for the two-player game with transferable utility, the set of subgame perfect payoffs contains only the unique SSP

payoff. We also establish restrictions on the discount factor which insure the uniqueness of the SP payoff for any n-player game.

Section 7 concludes with some observations on the implications of assuming that utility is transferable.

2. The Game

Let $\sigma = (\sigma_0, \sigma_1, \sigma_2, \dots)$ denote a Markov process realizing values in a Borel subset of a complete separable metric space, S . Let $K = \{1, \dots, k\}$ denote a set of players, where $k \geq 1$. We refer to an element $s \in S$ as a **state**, and an element $i \in K$ as a **player**. For $t = 0, 1, 2, \dots$, $\sigma^t = (\sigma_0, \dots, \sigma_t)$ denotes the t-period state-history with typical realization (s_0, \dots, s_t) .

A **stochastic sequential bargaining game** for K may be indexed by (c, ρ, β) , where, for each state s , $c(s)$ is a nonnegative real number representing the **cake** to be divided among the players if they agree in that state, $\rho(s)$ is a permutation on K with $\rho_i(s)$ denoting the identity of the player who makes the i^{th} move in that state, and β is a positive real number representing the factor by which utility attained in the next period is discounted. We suppose that $\beta < 1$, and that there is a finite \bar{c} such that $c(s) \leq \bar{c}$ for all $s \in S$. We also suppose that c and ρ are measurable functions of S .

To describe the game, it is useful to introduce some additional notation. Let \mathbb{R}^1 denote the real line and \mathbb{R}^k denote k-dimensional Euclidean space. For $x, y \in \mathbb{R}^k$, let $x \leq y$ denote $x_i \leq y_i$, $i = 1, \dots, k$, and let $\mathbb{R}_+^k = \{x \in \mathbb{R}^k: x \geq 0\}$. The k-vector $(0, \dots, 0)$ will be denoted by 0 . For any $s \in S$, let $C(s) = \{x \in \mathbb{R}_+^k: \sum_{i=1}^k x_i \leq c(s)\}$ denote the set of feasible utility

vectors to be allocated in state s . For an **allocation** $x \in C(s)$, x_i is the amount of cake awarded to player i . Let $\kappa(s) = \rho_1(s)$ denote the **proposer** in state s .

The game is played as follows. Upon the realization of a state s , $\kappa(s)$ 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. Each player responds in turn (in the order prescribed by ρ) until either some player has rejected the offer or all players have accepted it. If no proposal is offered and accepted by all players, a new state s' is realized in the next period according to the Markov process σ . The procedure is then repeated except that the order of moves is determined by $\rho(s')$ and the proposal must lie in the set $C(s')$. This process continues until an allocation is proposed and accepted.

An **outcome** (η, τ) may be defined as stopping time τ ,² and a k -random variable, η , measurable with respect to σ^τ , which satisfies $\eta \in C(\sigma_\tau)$ if $\tau < \infty$, and $\eta = 0$, otherwise. Given a realization of σ , τ denotes the number of periods which elapse before a proposal is accepted, and η denotes the allocation at the realized state σ_τ . An outcome implies a von-Neumann Morgenstern **payoff** to player i , $E[\beta^\tau \eta_i | \sigma_0 = s]$, for the game starting in state s . Our assumptions on c and β imply that the payoff is 0 if no proposal is ever accepted.

An outcome (η, τ) is **stationary** if there is a measurable function $\mu: S \rightarrow R^k$, and a measurable subset $S^\mu \subset S$ such that (i) $\sigma_t \notin S^\mu$, $t = 0, \dots, \tau-1$,

² A random variable τ is a stopping time on σ if τ is a function of σ taking values in $(0, 1, 2, \dots)$, and $\tau(\sigma)$ depends only on σ^τ .

(ii) $\sigma_r \in S^\mu$, and (iii) $\eta = \mu(\sigma_r)$. That is, upon reaching state s , proposal $\mu(s)$ is accepted if and only if $s \in S^\mu$. We typically represent a stationary outcome by (μ, S^μ) and its associated vector of payoffs at s by $v^\mu(s)$. S^μ is referred to as its **agreement set**. We may assume that $\mu(s) = 0$ for $s \in S - S^\mu$.

Since the remaining concepts are standard, we present them informally. A **history** is a specification of a finite sequence of realized states and the actions taken at each state in the sequence up to that point. A **strategy** for player i specifies a feasible action at every history at which he must act. A **strategy profile** is a measurable k -tuple of strategies,³ one for each player.

At any history, a strategy profile induces an outcome and hence a payoff for each player. A strategy profile is **subgame perfect** (SP) if, at every history, it is a best response to itself for the player who moves. We refer to the outcome and payoff functions induced by a subgame perfect strategy profile as an SP outcome and SP payoff respectively.

A strategy profile is **stationary** if the actions prescribed at any history depend only on the current state and current offer. A **stationary subgame perfect** (SSP) outcome and payoff is an outcome and payoff generated by subgame perfect strategy profile which is stationary. For any SSP payoff $v: S \rightarrow R^k$, we refer to $\sum_{i \in K} v_i: S \rightarrow R^1$ as its **total SSP payoff**.

We restrict attention to pure strategies only for notational convenience. As will be apparent from the analysis, allowing mixed strategies does not change the set of SP or SPP payoffs. Furthermore, since the state

³ That is, each strategy is measurable with respect to the Borel subsets of the product topology on the space of histories.

space need not be finite, the game itself need not be stationary in any sense. Time dependent strategies or even strategies which depend only on the realization of all past states may be incorporated into the set of stationary strategies simply by modifying the definition of the state space and the corresponding transition function. For instance, if we suppose that the transition probability has a distinct support for each state, then the set of outcomes and stationary outcomes essentially coincide. A stationary strategy profile is then simply a strategy profile which is independent of the previous actions of the players. Similarly, a finite horizon may be incorporated into the model by imposing restrictions on the transition functions such that some absorbing state s with $c(s) = 0$ is always realized by some finite period T .

3. **An Example**

Before proceeding to the general analysis, we present a simple example to illustrate the kind of results one can obtain when the cake size and proposer follow non trivial stochastic processes. To keep the analysis brief and to focus on the central issues, we freely exploit the more general results to be established in the next section.

There are two players, 1 and 2, and two possible cake sizes, c_1, c_2 , with $c_1 > c_2 > 0$. Given a current cake size c_i , cake size c_j is realized in the next period with probability q_{ij} . If cake size c_i is realized, then, with probability π_{ij} , player j is the proposer. Given these assumptions, we may suppose that $S = \{(i,j): i,j = 1,2\}$, where, given state (i,j) , c_i is the cake size, player j is the proposer, and $q_{ih}\pi_{h\ell}$ is the probability that state (h,ℓ) is realized in the next period.

As we show in Sections 4 and 5, this game has a unique SP payoff which is also SSP. It is sufficient, therefore, to determine the unique SSP payoff,

$v: S \rightarrow R^2$. As we establish in Theorem 1 below, any SSP payoff is the fixed point of a functional equation relating the payoff at the current state to its discounted expected value in the next period.

Consider the SSP payoff to player j when the current cake size is c_i and the other player j' is the proposer. If no proposal is agreed upon in the current period, then the payoff to player j is the expected value of his payoff in the next period. Therefore, it is a best response for player j to accept any proposal which yields him at least this much utility, and, consequently, it cannot be a best response for player j' to offer him more. We conclude that, regardless of whether player j' decides to make an acceptable offer or pass, player j 's SSP payoff in state (i, j') is

$$(1) \quad v_j(i, j') = \beta \sum_{(h, \ell) \in S} q_{ih} \pi_{h\ell} v_j(h, \ell).$$

Next consider the SSP payoff to player j when the current cake size is c_i and he is the proposer. He must decide between making an acceptable offer or passing. If he passes, he earns his expected future discounted SSP payoff. On the other hand, the other player j' accepts any proposal which gives him more than $v_{j'}(i, j)$. Therefore, the best response of player j must be to offer player j' her reservation value, $v_{j'}(i, j)$, whenever the remaining surplus exceeds her expected discounted future payoff from passing. We conclude that his SSP payoff in state (i, j) must satisfy

$$(2) \quad v_j(i, j) = \max\left(c_i - \beta \sum_{(h, \ell) \in S} q_{ih} \pi_{h\ell} v_{j'}(h, \ell), \beta \sum_{(h, \ell) \in S} q_{ih} \pi_{h\ell} v_j(h, \ell)\right).$$

Let $w(i,j) = \sum_{h=1}^2 v_h(i,j)$ denote the SSP total payoff at state (i,j) . Then summing over equations (1) and (2), we obtain a functional equation for the total payoff.

$$(3) \quad w(i,j) = \max\{c_i, \beta \sum_{(h,\ell) \in S} q_{ih} \pi_{h\ell} w(h,\ell)\}, \quad i,j = 1,2.$$

Notice that $w(i,j)$ depends on j only to the extent that it affects the transition probability, $q_{ih} \pi_{h\ell}$. Since, in this example the transition probability is independent of the proposer, it follows that the total payoff depends only on the current cake size, and we may write $w(i,j) = w_i$, $i,j = 1,2$. Equation (3) may then be expressed as

$$(3') \quad w_i = \max\{c_i, \beta \sum_{j=1}^2 q_{ij} w_j\}, \quad i = 1,2.$$

Since $c_1 > c_2$, the reader may verify that the unique solution to this system of equations is

$$(4) \quad w_1 = c_1,$$

and

$$(5) \quad w_2 = \max\{c_2, \beta c_1 q_{21} / (1 - \beta q_{22})\}.$$

In an SSP equilibrium, players always consume the entire cake in any agreement. Consequently, there is no benefit from either passing or rejecting an SSP offer whenever the total payoff is equal to the current cake size.

Therefore, there is always an SSP strategy profile which implies agreement in any state (i,j) at which $w_1 = c_1$. In particular, the set of states in which agreement occurs does not depend on the identity of the current proposer.

Returning to the individual payoffs, we may use the definition of w and relation (1) to rewrite the expression for $v_j(i,j)$ in (2) solely in terms of v_j , w , and the transition probabilities,

$$(2') \quad v_j(i,j) = \max\{c_1 - \beta(q_{11}w_1 + q_{12}w_2), 0\} + v_j(i,j'), \quad j' \neq j.$$

There are two cases to consider.

Case I: $w_2 = \beta c_1 q_{21} / (1 - \beta q_{22}) > c_2$, so that agreement occurs only when the cake size is c_1 . In this case, $c_2 < w_2 = \beta[q_{21}w_1 + q_{22}w_2]$ and $c_1 = w_1 > \beta[q_{11}w_1 + q_{12}w_2]$, so that equation (2') reduces to

$$v_j(1,j) = c_1 - \beta(q_{11}w_1 + q_{12}w_2) + v_j(1,j'), \quad j = 1,2,$$

and

$$v_j(2,j) = v_j(2,j'), \quad j = 1,2, \quad j' \neq j.$$

Case II: $w_2 = c_2 > \beta c_1 q_{21} / (1 - \beta q_{22})$. In this case, agreement occurs in all states and for $i = 1,2, j = 1,2$,

$$v_j(i,j) = c_1 - \beta(q_{11}c_1 + q_{12}c_2) + v_j(i,j'), \quad j = 1,2, \quad j' \neq j.$$

We conclude that if no agreement is reached at the current state, then there is no advantage to a player from being the proposer. If agreement is reached at the current state, then the gain to a player from being the proposer is precisely the difference between the current cake size and expected discounted value of the future cake which will be agreed upon.

Although the net gain from proposing depends only on the distribution of cake sizes for the agreement states, the total payoff to a player, even when he is responding does depend on the future identity of the proposer for different cake sizes. To illustrate, suppose that $\pi_{ij} = \alpha_j$, $i, j = 1, 2$, so that a given player has the same probability of being the proposer in each period, independent of the cake size, and $q_{ij} = 1/2$, $i, j = 1, 2$, so that either cake size has an equal probability of being chosen in each period. Then, upon solving equations (1) and (2'), we obtain as the equilibrium payoff to the responder j in any agreement state (i, j') ,

$$v_j(i, j') = \beta \alpha_j (w_1 + w_2) / 2.$$

Not only is the payoff to player j independent of the current cake size, but it converges to 0 as α_j goes to 0.

This point can be illustrated even more sharply if we suppose that player i moves only when the cake is size c_1 , $i = 1, 2$. Suppose that $c_2 > \beta c_1 q_{21} / (1 - \beta q_{22})$ so agreement is always reached in every state. In this case, the solution to (1) and (2') yields

$$v_1(2, 2) = A(c_1 - Bc_2) / (1 - AB),$$

and

$$v_2(1,1) = B(c_2 - Ac_1)/(1 - AB),$$

where $A = \beta q_{21}/(1-\beta q_{22})$ and $B = \beta q_{12}/(1-\beta q_{11})$. Notice, however, that as c_2 converges to $\beta c_1 q_{21}/(1-\beta q_{22})$, $v_2(1,1)$ goes to 0. In fact, if $c_2 < \beta c_1 q_{21}/(1-\beta q_{22})$ so that agreement occurs only when the cake is size c_1 , the SSP payoff for player 2 is 0. The reason is that he never obtains any surplus so that his reservation value is consequently 0 in every state. This result is particularly striking when we observe that for any difference between c_1 and c_2 , there is a β sufficiently close to 1 such that agreement occurs only at cake size c_1 , resulting in an equilibrium payoff of 0 for player 2.

These examples contain a number of features which may be established for any game satisfying the assumptions of Section 2. Equations (1) and (2) from which all of the results of the example are derived are a special case of a general characterization of any SSP payoff for any finite number of players. It is also a general property of any game with transferable utility that the total payoff maximizes the expected discounted value of the cake size agreed upon. Consequently, the set of agreement states depends only on the process of cake sizes, independently of the process which determines the proposers. The relation between the advantage to proposing and net gain to an agreement in the current period is also a general feature of any game with transferable utility.

4. Stationary Subgame Perfect Payoffs

In this section, we focus on the SSP payoffs. The key feature of an SSP equilibrium is that there is no role for threats to prevent the proposer

from fully exploiting his monopoly power, thereby inducing an inefficient allocation. This property also permits a sharp characterization of the SSP payoffs in terms of the fixed point of an operator on the set of payoff functions. Combined with our assumption of transferable utility, this characterization allows an easy proof of the uniqueness of the SSP payoff even for games with more than two players. We are also able to establish strong results with respect to the advantage to proposing and the conditions under which agreement is delayed. We conclude the section by establishing some conditions under which agreement follows a reservation rule on the size of the cake.

One may wish to study the SSP payoff as a focal point of the set of SP payoffs. For example, Baron and Kalai (1990) show that SSP equilibria minimize the number of automaton states required to implement an SP equilibrium. However, the SSP payoff is also worth studying because of its relation to the set of SP payoffs. As we report in Section 5, for two-player games, the uniqueness of SSP payoff implies the uniqueness of the SP payoff. For more than two players, the SP payoffs generally exhibit extreme indeterminacy except for very low discount factors. Nevertheless, we may still formulate a characterization of the SP payoffs which is based on the same operator used to characterize the SSP payoffs.

4.1. A Characterization of SSP payoffs: Existence and Uniqueness

We begin by establishing a general characterization of SSP payoffs as the set of fixed points of a certain operator on the set of k -dimensional payoff functions on S . The operator is based on the observation that, if agreement is reached in any period, the proposer may extract any surplus over

what the players obtain by delaying agreement until the next period. We then exploit this characterization to establish a similar characterization for the set of total SSP payoffs which depends only on the process of cake sizes. This observation leads in turn to an alternative characterization of the SSP payoff for each player as the fixed point of an operator which depends only on the process of cake sizes and the set of states in which he is the proposer. The existence of a unique solution is then established by demonstrating that the operators whose fixed points characterize the total SSP payoff and individual SSP payoffs satisfy the contraction property.

For any n -dimensional Euclidean space R^n , let F^n denote the set of bounded measurable functions on S taking values in R^n . For $f^1, f^2 \in F^n$, let $f^1 \leq f^2$ denote $f^1(s) \leq f^2(s)$, $s \in S$. For $x \in R^n$, let $\|x\| = \sup\{|x_i| : i = 1, \dots, n\}$, and for any $f \in F^n$, let $\|f\| = \sup\{\|f(s)\| : s \in S\}$. An operator $Q: F^n \rightarrow F^n$ satisfies the **contraction** property if there is a $\delta < 1$ such that $f^1, f^2 \in F^n$ implies $\|Q(f^1) - Q(f^2)\| < \delta \|f^1 - f^2\|$. For any operator $Q: F^n \rightarrow F^n$ which satisfies the contraction property, there is a unique $f \in F^n$ such that $Q(f) = f$ (see e.g. Blackwell (1965), p. 232).

For any stationary outcome (μ, S^μ) and $f \in F^k$, define $V^\mu(f)(s) = \mu(s)$ if $s \in S^\mu$, and $V^\mu(f)(s) = \beta E[f(\sigma_1) | \sigma_0 = s]$, otherwise. If $f(s')$ defines the undiscounted payoff upon reaching state s' in the next period, then $V^\mu(f)(s)$ defines the expected payoff from outcome (μ, S^μ) at state s . Note that, for any $s \in S$, $V^\mu(v^\mu)(s) = v^\mu(s)$. Furthermore, $V^\mu: F^k \rightarrow F^k$ satisfies the contraction property since, for $f^1, f^2 \in F^k$, and $s \in S$, either $\|V^\mu(f^1)(s) - V^\mu(f^2)(s)\| = 0$, or $\|V^\mu(f^1)(s) - V^\mu(f^2)(s)\| = \beta \|E[f^1(\sigma_1) - f^2(\sigma_1) | \sigma_0 = s]\| \leq \beta E[\|f^1(\sigma_1) - f^2(\sigma_1)\| | \sigma_0 = s] \leq \beta \|f^1 - f^2\|$. Therefore, v^μ is the unique fixed point of V^μ .

With these preliminaries we turn to the characterization of the set of SSP payoffs. For $f \in F^k$, define $A_i(f)(s) = \max\{c(s) - \beta \sum_{j \in K} E[f_j(\sigma_1) | \sigma_0 = s], 0\} + \beta E[f_i(\sigma_1) | \sigma_0 = s]$ for $i = \kappa(s)$, and $A_j(f)(s) = \beta E[f_j(\sigma_1) | \sigma_0 = s]$, $j \neq \kappa(s)$. Then $A(f)(s)$ denotes the equilibrium payoff of the ultimatum game in state s when the payoff to disagreement is $\beta E[f(\sigma_1) | \sigma_0 = s]$.

Theorem 1: f is an SSP payoff if and only if $A(f) = f$.

Proof: (\Rightarrow) Suppose f is SSP. Fix $s \in S$ and let $i = \kappa(s)$. If no proposal is accepted in state s , then the SSP payoff is $\beta E[f(\sigma_1) | \sigma_0 = s]$. Therefore, it is a best response for player $j \neq i$ to accept a proposal $x \in C(s)$ if and only if $x_j \geq \beta E[f_j(\sigma_1) | \sigma_0 = s]$. Consequently, $A(f)(s)$ is the only acceptable proposal which is best response for player i . Therefore, in any SSP outcome with payoff f , no proposal is offered and accepted if $c(s) < \beta \sum_{j \in K} E[f_j(\sigma_1) | \sigma_0 = s]$, and proposal $A(f)(s)$ must be accepted if $c(s) > \beta \sum_{j \in K} E[f_j(\sigma_1) | \sigma_0 = s]$. Otherwise, there is no best response for player i . It then follows by definition that $A(f) = f$.

(\Leftarrow) Suppose $A(f) = f$. We will construct an SSP outcome (μ, S^μ) . Let $S^\mu = \{s \in S: f(s) \in C(s)\}$, and let $\mu(s) = f(s)$ for $s \in S^\mu$, and $\mu(s) = 0$, otherwise. By construction, $V^\mu(f) = A(f) = f$, and therefore $f = v^\mu$. Now consider the stationary strategy in which, upon reaching any state s , player $\kappa(s)$ proposes $\mu(s)$, and any player $j \neq \kappa(s)$ accepts any proposal which yields him a payoff no less than $v_j^\mu(s)$. The theorem is proved by noting that, given the future return function v^μ , any player j who unilaterally defects from the prescribed strategy in state s earns a return no greater than $v_j^\mu(s)$. Q.E.D.

Recall that a function $w: S \rightarrow R^1$ is a total SSP payoff if there is an SSP payoff, v , such that $w = \sum_{i \in K} v_i$. Theorem 1 leads directly to a characterization of the set of total SSP payoffs. For any real valued bounded function h on S , let $H(h)(s) = \max(c(s), \beta E[h(\sigma_1) | \sigma_0 = s])$, $s \in S$.

Theorem 2: If w is an SSP total payoff, then $H(w) = w$.

Proof: Suppose v is a SSP payoff. Let $w = \sum_{i \in K} v_i$. Then, since Theorem 1 implies that $A(v) = v$, it follows by definition that $H(w) = \sum_{i \in K} A_i(v) = \sum_{i \in K} v_i = w$. Q.E.D.

One striking implication of Theorem 2 is that the SSP total payoff, and hence the set of states in which agreement occurs, does not depend on which player is the proposer in any state.

Corollary 1: The total payoff depends only on $\{c(\sigma_t): t = 0, 1, 2, \dots\}$.

In fact, H defines the "maximization" operator for the single agent problem of deciding when to consume a stochastic cake in which the maximum size of the cake declines exponentially with time. We conclude that the SSP total payoff maximizes the expected surplus allocated among the players. The order in which players move affects how the cake is allocated, but not the state in which it is allocated. Notice that this result requires no additional restrictions on the size of the cake or discount factors, the order of moves, nor the Markov process σ .

Combining Theorems 1 and 2, we may provide an alternative characterization of the SSP payoffs which allows the SSP payoff to each player to be calculated independently. Given a total payoff function $w \in F^1$, we define an operator $B_i(\cdot; w): F^1 \rightarrow F^1$ for each player i as follows. For $i = \kappa(s)$, let $B_i(f_i; w)(s) = \max\{c(s) - \beta E[w(\sigma_1) | \sigma_0 = s], 0\} + \beta E[f_i(\sigma_1) | \sigma_0 = s]$, and, for $i \neq \kappa(s)$, let $B_i(f_i; w)(s) = \beta E[f_i(\sigma_1) | \sigma_0 = s]$. Although the payoffs of the other players do not explicitly appear in $B_i(f_i; w)$, they are implicit whenever player i is the proposer, since $w - v_i = \sum_{j \neq i} v_j$.

For $f \in F^k$, let $B(f; w) = (B_1(f_1; w), \dots, B_k(f_k; w))$. Then $B(\cdot; w): F^k \rightarrow F^k$ defines an operator on F^k .

Lemma 1: Suppose $H(w) = w$. Then v is an SSP payoff if and only if $B(v; w) = v$.

Proof: (\Rightarrow) If v is an SSP payoff, then Theorem 1 and Theorem 2 imply $A(v) = v$ and $H(\sum_{i \in K} v_i) = \sum_{i \in K} v_i$. Letting $w = \sum_{i \in K} v_i$, it then follows from the definition of B , that $B(v; w)(s) = A(v)(s) = v(s)$, $s \in S$.

(\Leftarrow) Suppose $H(w) = w$ and $B(v; w) = v$. Then $\sum_{i \in K} v_i = \sum_{i \in K} B_i(v_i; w) = \max\{c(s) - \beta E[w(\sigma_1) | \sigma_0 = s], 0\} + \beta \sum_{i \in K} E[v_i(\sigma_1) | \sigma_0 = s]$. Now consider the operator \hat{H} on F^1 defined by $\hat{H}(h)(s) = \max\{c(s) - \beta E[w(\sigma_1) | \sigma_0 = s], 0\} + \beta E[h(\sigma_1) | \sigma_0 = s]$. Note that \hat{H} is a contraction since, for $h^1, h^2 \in F^1$, $|\hat{H}(h^1)(s) - \hat{H}(h^2)(s)| = \beta |E[h^1(\sigma_1) | \sigma_0 = s] - E[h^2(\sigma_1) | \sigma_0 = s]| \leq \beta E[|h^1(\sigma_1) - h^2(\sigma_1)| | \sigma_0 = s] \leq \beta \|h^1 - h^2\|$. Therefore, $\sum_{i \in K} v_i$ is the unique fixed point of \hat{H} . But, by definition, $\hat{H}(w)(s) = \max\{c(s) - \beta E[w(\sigma_1) | \sigma_0 = s], 0\} + \beta E[w(\sigma_1) | \sigma_0 = s] = H(w)(s) = w(s)$, $s \in S$. Therefore, $w = \hat{H}(w)$, and hence $w = \sum_{i \in K} v_i$. Substituting for w in the definition of B , we obtain $v = B(v; w) = A(v)$. The result then follows from Theorem 1. Q.E.D.

To establish the existence of a unique SPP payoff, it is sufficient to show that H possesses a unique fixed point w for which $B(\cdot;w)$ also has a unique fixed point.

Theorem 3: There is a unique SSP payoff.

Proof: To show H satisfies the contraction property, suppose $h^1, h^2 \in F^1$, and consider any $s \in S$. Without loss of generality, we may assume that $H(h^1)(s) \geq H(h^2)(s)$. Then, there are two cases to consider. If $\beta E[h^1(\sigma_1) | \sigma_0=s] > c(s)$, then $|H(h^1)(s) - H(h^2)(s)| = \beta E[h^1(\sigma_1) | \sigma_0=s] - \max(c(s), \beta E[h^2(\sigma_1) | \sigma_0=s]) \leq \beta E[h^1(\sigma_1) | \sigma_0=s] - \beta E[h^2(\sigma_1) | \sigma_0=s] \leq \beta E[|h^1(\sigma_1) - h^2(\sigma_1)| | \sigma_0=s] \leq \beta \|h^1 - h^2\|$. Otherwise, $c(s) = H(h^1)(s) \geq H(h^2)(s) = \max(c(s), \beta E[h^2(\sigma_1) | \sigma_0=s])$ implies $H(h^2)(s) = c(s)$, in which case $|H(h^1)(s) - H(h^2)(s)| = 0$. This establishes that H possesses the contraction property and hence that there is a unique w such that $w = H(w)$.

To show that $B_i(\cdot;w)$ satisfies the contraction property, we note that, for $h^1, h^2 \in F^1$, $|B_i(h^1)(s) - B_i(h^2)(s)| = \beta |E[h^1(\sigma_1) | \sigma_0=s] - E[h^2(\sigma_1) | \sigma_0=s]| \leq \beta E[|h^1(\sigma_1) - h^2(\sigma_1)| | \sigma_0=s] \leq \beta \|h^1 - h^2\|$. We conclude that there is a unique v such that $B(v;w) = v$. The result then follows from Lemma 1. Q.E.D.

4.2 Advantage to Proposing

Given that the set of agreement states does not depend on which player proposes, the only role for the proposer process is to determine the allocation of the cake in those states where agreement is reached. It is clear from the definition of B_i that, given any payoff function for future

periods, a player can never be hurt by being a proposer since he always has the option of passing and receiving the same expected discounted payoff as he would if he were a responder. However, we may exploit Lemma 1 to establish a much stronger result. If a game is altered so that the set of states in which some player proposes is increased, not only does he benefit in those states in which he is now the proposer, but his SSP payoff does not decrease in any other state either. However, there is only an advantage to proposing in those states in which agreement occurs. If a player is never a proposer in an agreement state, his SSP payoff is 0 in all states.

Consider two games $G^j = (c, \rho^j, \beta)$, $j = 1, 2$, which differ only in who makes the proposal in any state, and let v^j , $j = 1, 2$, denote the corresponding SSP payoffs.

Theorem 4: (i) If, for all $s \in S$, $\kappa^1(s) = i$ implies $\kappa^2(s) = i$, then $v_i^2 \geq v_i^1$. (ii) If, for all $s \in S$ such that $c(s) \geq \beta E[w(\sigma_1) | \sigma_0 = s]$, $\kappa^1(s) = i$ if and only if $\kappa^2(s) = i$, then $v_i^1 = v_i^2$.

Proof: Let B^j denote the B operator associated with game G^j , $j = 1, 2$, and let w denote the SSP total payoff for both games.

(i) Let $h^0 = v_i^1$ and consider the sequence (h^t) , defined by $h^t = B_i^2(h^{t-1})$, $t = 1, 2, 3, \dots$. Then, since $\kappa^2(s) = i$ whenever $\kappa^1(s) = i$, the definition of B implies that $h^1 = B_i^2(h^0; w) \geq B_i^1(h^0; w) = B_i^1(v_i^1; w) = v_i^1 = h^0$. Proceeding by induction, it then follows from the definition of B_i^2 that $h^t = B_i^2(h^{t-1}; w) \geq B_i^2(h^{t-2}; w) = h^{t-1}$, $t = 1, 2, 3, \dots$. But, since $B^2(\cdot; w)$ is a contraction, it follows that $h^t \uparrow v_i^2 = B_i^2(v_i^2; w)$, and hence that $v_i^1 \leq v_i^2$.

(ii) If $c(s) \geq \beta E[w(\sigma_1) | \sigma_0 = s]$ implies $\kappa^1(s) = i$ if and only if $\kappa^2(s) = i$, $s \in S$, then, it follows from the definition of B , that $B^1 = B^2$. The result then follows from Lemma 1 and Theorem 3. Q.E.D.

Condition (i) states that if the game is altered so that player i becomes a proposer in a new state, then his new SSP payoff does not decrease in any state. The reason is that by increasing his payoff in the new state, the player raises his reservation value in any state that immediately precedes it with positive probability. Furthermore, since the total payoff is fixed, the payoff to the remaining players can only decrease and hence their reservation value in any preceding states can only decrease. Consequently, the feedback effect always benefits the new proposer at the expense of the other players.

Condition (ii) implies that the only states in which there is a benefit to proposing are those in which agreement occurs. If agreement does not occur, the payoff associated with that state is just the expected future payoffs of those states in which agreement does occur. Similarly, if a player is a responder, he will never be offered more than the expected value of his future payoffs. Since future payoffs are discounted, it follows that if he ever earns a positive payoff, his highest payoffs must be in states where he makes an acceptable offer. If no such state exists, his SSP payoff must be zero.

4.3 Necessary and Sufficient Conditions for Agreement

In their discussion of nonstochastic sequential bargaining models, Osborne and Rubinstein (1990) emphasize that in an environment of complete

information, agreement should generally occur immediately as long as the subgame perfect equilibrium is unique. When the cake is stochastic, the result needs to be modified. What is preserved is not agreement, but efficiency.

As we report in Section 5, if the SP payoff is unique, it must be SSP. But in any SSP equilibrium, the total payoff maximizes the expected discounted value of the cake. This implies immediately that, if the cake size is constant, agreement must occur in every state. However, at any state in which there is an outcome for which future expected discounted cake size exceeds the current cake size, agreement will be delayed. Our next result demonstrates that such an outcome can be found for some state if and only if the process of discounted cake sizes is not a supermartingale.

Theorem 5: If v is the SSP payoff, then there is an SSP outcome for which agreement occurs in every state if and only if, $\beta E[c(\sigma_1) | \sigma_0 = s] \leq c(s)$, $s \in S$.

Proof: (\Rightarrow) Suppose that $\beta E[c(\sigma_1) | \sigma_0 = s] > c(s)$ for some $s \in S$. Then if w is the total payoff implied by SSP outcome (μ, S^μ) , Theorem 2 implies $w(s) = H(w)(s) \geq \beta E[w(\sigma_1) | \sigma_0 = s] \geq \beta E[c(\sigma_1) | \sigma_0 = s] > c(s)$, and hence that $s \notin S^\mu$.

(\Leftarrow) Suppose that $\beta E[c(\sigma_1) | \sigma_0 = s] \leq c(s)$ for all $s \in S$. Then σ a Markov process implies $E[\beta^{\tau+1} c(\sigma_{\tau+1}) | \sigma^\tau] \leq c(\sigma_\tau)$. Given Theorem 2, it suffices to show that $c = H(c)$. Suppose (μ, S^μ) is the SSP allocation and let (η, τ) denote the random allocation and stopping time associated with (μ, S^μ) . Then, $\sum_{i \in K} \eta_i \leq c(\sigma_\tau)$. Consider any state $s \notin S^\mu$. Then, $w^\mu(s) = E[\beta^\tau \sum_{i \in K} \eta_i | \sigma_0 = s] \leq E[\beta^\tau c(\sigma_\tau) | \sigma_0 = s]$. But, since $\beta E[c(\sigma_1) | \sigma_0 = s] \leq c(s)$ for all $s \in S$, and σ is a Markov process, we have, $E[\beta^{t+1} c(\sigma_{t+1}) | \sigma^t] \leq \beta^t c(\sigma_t)$, $t = 0, 1, 2, \dots$. Therefore,

since $\beta^t c(\sigma_t) \rightarrow 0$ as $t \rightarrow \infty$, the optional stopping theorem (see e.g. Karlin and Taylor (1975), p. 267) implies $E[\beta^t c(\sigma_t) | \sigma_0 = s] \leq c(s)$, and hence that $w^\mu(s) \leq c(s)$. Therefore, (μ, S^μ) may be chosen so that $s \in S^\mu$. Q.E.D.

Since we take the cake size as a primitive of the model, it is not immediately clear from Theorem 5 which bargaining contexts are most likely to exhibit protracted negotiations. To address this question, a brief discussion of the fundamentals underlying the cake process is in order.

The cake size is the sum of the payoffs to the k players attainable by the best agreement that can be reached in the current state. It does not depend on when the terms of the agreement are to be implemented or when the discounted returns accrue. It only requires that the players commit themselves to all future state contingent actions prescribed by the agreement. Consequently, if all conceivable state contingent actions can be agreed upon in the current state and agents evaluate the flow of benefits according to a time separable von Neumann-Morgenstern utility, the current cake size must exceed the expected discounted value of any alternative outcome. Any efficiency gain from waiting must derive from limitations on the set of feasible agreements in the current state.

Suppose, for example, that two players are bargaining over a car. The current owner derives a benefit of \$100 per period until period T after which he has no use for the car, whereas, the buyer would immediately derive a perpetual benefit of \$50 per period. Clearly, the car should not be physically exchanged until period T in any efficient agreement. Nevertheless, the agreement to exchange the car in period T could be made immediately. Given a common discount factor of β , the cake size in period $t \leq T$ is just the

discounted value of cake size in period T, $c(t) = \beta^{T-t}c(T) = 50 \beta^{T-t}/(1-\beta)$.

If, however, it is not feasible to agree upon a future exchange, then from period 0 to period T, the discounted cake size actually grows since an early exchange implies an efficiency loss of \$50 per period in the first T periods. Specifically, for $t \leq T$, $c(t) = \beta^{T-t}c(T) - 50(1-\beta^{T-t})/(1-\beta) < \beta^{T-t}c(T)$.

In general, the more complicated the environment, the more difficult it may be to specify the terms of an efficient state contingent agreement. This constraint may be particularly severe when the objects to be exchanged can be described only by a high dimensional vector of attributes and the process of formulating offers and counter offers takes some time, such as in negotiations over a labor contract or some political agreement. Merlo (1992), in his analysis of government formation in a parliamentary system, uses a special case of this model to explain the empirical relation between the duration of negotiations among the members of the ruling coalition and the state of the economy.

4.4 Sufficient Conditions for a Reservation Rule for Agreement

In this section we establish conditions on $c(\sigma_t)$ which imply that agreement follows a reservation rule on the size of the cake. For any set D, let l_D denote its indicator function.

$$(D) \quad c(s') \geq c(s'') \text{ implies (i) } E[l_{(-a,b)}(c(\sigma_1)) | \sigma_0=s'] \leq E[l_{(-a,b)}(c(\sigma_1)) | \sigma_0=s''], \quad b \in R^1, \text{ (ii) } E[c(\sigma_1) | \sigma_0=s'] - E[c(\sigma_1) | \sigma_0=s''] \leq \beta^{-1}[c(s') - c(s'')].$$

Condition (D) imposes two restrictions on $c(\sigma_t)$. The conditional distribution of the next period cake size must satisfy first order stochastic

dominance with respect to the current cake size. However, the average next period cake size cannot increase at a rate greater than β^{-1} with respect to the current cake size.

Theorem 6: If condition (D) is satisfied, then there is an SSP outcome such that agreement occurs in state s if and only if $c(s) \geq c^*$.

Proof: Suppose v is the SSP payoff and w is the total SSP payoff. We will show that $|w(s') - w(s'')| \leq |c(s') - c(s'')|$ for $s', s'' \in S$. Then, since the definition of H and Theorem 2 imply $c \leq H(w) = w$, it follows that there is a $c^* \geq 0$ such that $w(s) = c(s)$ if and only if $c(s) \geq c^*$. Letting $S^\mu = \{s \in S: c(s) \geq c^*\}$ and $\mu(s) = v(s)$ for $s \in S^\mu$ establishes the result.

Let H^t denote the t -fold application of H . Since H is a contraction with fixed point w , it follows that $H^t(c) \rightarrow w$ as $t \rightarrow \infty$. Therefore, to establish that $|w(s') - w(s'')| \leq |c(s') - c(s'')|$ for $s', s'' \in S$, it is sufficient to show that, for any real valued function h on S , $0 \leq h(s') - h(s'') \leq c(s') - c(s'')$ for $s', s'' \in S$, implies $0 \leq H(h)(s') - H(h)(s'') \leq c(s') - c(s'')$ for $s', s'' \in S$.

Fix a function h such that $c(s') \geq c(s'')$ implies $0 \leq h(s') - h(s'') \leq c(s') - c(s'')$ for $s', s'' \in S$. Note first that this restriction implies a nondecreasing function $\psi_1: \mathbb{R}^1 \rightarrow \mathbb{R}^1$ such that $h(s) = \psi_1(c(s))$, $s \in S$, and a nonincreasing function ψ_2 such that $h(s) - c(s) = \psi_2(c(s))$, $s \in S$. Suppose $c(s') > c(s'')$.

We show first that $H(h)(s') - H(h)(s'') \leq c(s') - c(s'')$. By definition, $H(h)(s) = \max(c(s), \beta E[h(\sigma_1) | \sigma_0 = s])$, $s \in S$. Therefore, if $H(h)(s') = c(s')$, the result follows immediately. So suppose that $H(h)(s') >$

$c(s')$. Then the definition of H and Condition (D(ii)) imply $H(h)(s') - H(h)(s'') \leq \beta E[h(\sigma_1)|\sigma_0=s'] - \beta E[h(\sigma_1)|\sigma_0=s''] = \beta(E[h(\sigma_1)-c(\sigma_1)|\sigma_0=s'] - E[h(\sigma_1)-c(\sigma_1)|\sigma_0=s'']) + \beta(E[c(\sigma_1)|\sigma_0=s'] - E[c(\sigma_1)|\sigma_0=s'']) \leq \beta(E[h(\sigma_1)-c(\sigma_1)|\sigma_0=s'] - E[h(\sigma_1)-c(\sigma_1)|\sigma_0=s'']) + c(s') - c(s'')$.

For $b \in \mathbb{R}^1$, $s \in S$, let $P(b|s) = E[1_{(-\infty, b]}(h(\sigma_1)-c(\sigma_1))|\sigma_0=s] = E[1_{(-\infty, b]}(\psi_2(c(\sigma_1))|\sigma_0=s]$ denote the probability that $h-c$ is no greater than b in the next period, conditional on the current state s . Then $E[h(\sigma_1)-c(\sigma_1)|\sigma_0=s] = \int z dP(z|s)$, $s \in S$. Also, since ψ_2 is nonincreasing, condition (D(i)) implies that $P(\cdot|s') \geq P(\cdot|s'')$. Therefore, integrating by parts, we obtain, $H(h)(s') - H(h)(s'') \leq \beta(\int z dP(z|s') - \int z dP(z|s'')) + c(s') - c(s'') = \beta(\int P(z|s'') dz - \int P(z|s') dz) + c(s') - c(s'') \leq c(s') - c(s'')$.

The proof that $H(h)(s') \geq H(h)(s'')$ is similar. By definition, $H(h)(s) = \max\{c(s), \beta E[h(\sigma_1)|\sigma_0=s]\}$, $s \in S$. Therefore, if $H(h)(s'') = c(s'')$, the result follows immediately. So suppose that $H(h)(s'') > c(s'')$. Then the definition of H implies $H(h)(s') - H(h)(s'') \geq \beta(E[h(\sigma_1)|\sigma_0=s'] - E[h(\sigma_1)|\sigma_0=s''])$. For $b \in \mathbb{R}^1$, $s \in S$, let $Q(b|s) = E[1_{(-\infty, b]}(h(\sigma_1))|\sigma_0=s] = E[1_{(-\infty, b]}(\psi_1(c(\sigma_1))|\sigma_0=s]$ denote the probability that h is no greater than b in the next period, conditional on the current state s . Then, $E[h(\sigma_1)|\sigma_0=s] = \int z dQ(z|s)$, $s \in S$, and, since ψ_1 is nondecreasing, condition (D(i)) implies that $Q(\cdot|s') \leq Q(\cdot|s'')$. Therefore, integrating by parts, $H(h)(s') - H(h)(s'') \geq \beta(\int z dQ(z|s') - \int z dQ(z|s'')) = \beta(\int [Q(z|s'') - Q(z|s')] dz) \geq 0$. Q.E.D.

Condition (D) is by no means necessary for the SSP outcome to satisfy a reservation rule. For instance, if there are only two cake sizes, then any SSP outcome must satisfy a reservation rule since agreement will always occur at the largest cake. Nevertheless, if one wants to establish sufficient

conditions for a general model, something like condition (D) must be satisfied even for games with a relatively small number of states and degenerate transition probabilities. Consider a game with three cake sizes $\{1,2,3\}$ and let $q(j|i)$ denote the probability that cake size j is realized in the next period, given current cake size i . Without the stochastic dominance condition, we may admit $q(1|1) = 1$, $q(3|2) = 1$, and $q(1|3) = 1$, for which agreement occurs in states 1 and 3, but not in state 2 whenever the discount factor is greater than $2/3$. Even with the stochastic dominance condition, but without the restrictions on the rate of increase in the mean, we can admit $q(1|1) = 1$, $q(3|2) = 1$, and $q(3|3) = 1$, for which agreement again occurs in states 1 and 3, but not in state 2 whenever the discount factor is greater than $2/3$.

5. Subgame Perfect Payoffs

Once we drop the restriction that strategies must be stationary with respect to s , the characterization of the subgame perfect payoffs is complicated by the possibility that one or more players may employ nonstationary threats. Unless the SP payoff is unique, it is generally not possible to assign an equilibrium payoff to each state since the payoff at that state may depend on how it is reached. Nevertheless, since any threat must be subgame perfect, the most potent threats are the strategies which generate the extremal points of the SP payoffs. This observation, first exploited by Abreu (1988) for repeated games with discounting, leads to a characterization of the extremal SP payoffs as fixed points of an operator on $F^* \times F^*$ in which the best payoff for player i is computed against the worst payoff for all other players, and the worst payoff for player i is computed

against the best payoff of the other players. The question of the uniqueness of the SP payoff is then reduced to investigating conditions under which the best payoff must equal the worst payoff for each player.

Using this characterization for the extremal SP payoffs, we are able to establish three results. First, extremal SP payoffs always exist. Second, for a two-player game, the extremal SP payoffs are SSP payoffs. Therefore, if the SSP payoff is unique, then the SP payoff is unique. Third, for a general k -player game, the SP payoff is also unique if $\beta < 1/(k-1)$. All of these results are established under the more general conditions studied by Merlo and Wilson (1992). We will therefore confine ourselves to a brief introduction of the formal apparatus and some intuitive arguments behind the main results.

For $x \in R^k$, let $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$, and for $z \in R^1$, let $(z, x_{-i}) = (x_1, \dots, x_{i-1}, z, x_{i+1}, \dots, x_k)$, $i = 1, \dots, k$. Given $f, g \in F^k$, $s \in S$, $i \in K$, let $N_i(f, g)(s) = A_i(g_i, f_{-i})(s)$ denote the payoff to player i in state s for the ultimatum game in which the payoff to no agreement is $\beta E[(g_i, f_{-i})(\sigma_1) | \sigma_0 = s]$. If f_{-i} represents the best continuation payoffs for the players other than i , and g_i is his worst continuation payoff, then player i can do no worse than $N_i(f, g)(s)$ in state s .

Similarly, let $M_i(f, g)(s) = \max(c(s) - \sum_{j \neq i} g_j(s), \beta E[f_i(\sigma_1) | \sigma_0 = s])$ denote the maximum payoff to player i when an agreement must yield a payoff no less than $g_{-i}(s)$ to the other players and he can guarantee himself a return of $\beta E[f_i(\sigma_1) | \sigma_0 = s]$ if no agreement is reached. If g_{-i} represents the minimum payoffs to the players other than i in the current state, and f_i is his best continuation payoff, then player i can do no better than $M_i(f, g)$.

Note that when computing N_i , the reservation value of the other players is the expected discounted value of their maximum payoff in the next

period. In contrast, M_i is computed against a reservation value determined by the worst undiscounted payoff to other players in the current state.

Let $M \times N$ denote the operator on $F^k \times F^k$, defined by $M \times N(f, g) = (M(f, g), N(f, g))$. By exploiting the monotonicity of the M and N operators, Merlo and Wilson (1992) establish the following result.

Lemma 2: There is a pair $m^*, n^* \in F^k$, such that (i) $m^* \geq n^*$, (ii) $M \times N(m^*, n^*) = (m^*, n^*)$, and (iii) $M \times N(f, g) = (f, g)$ implies $n^* \leq f, g \leq m^*$.

A characterization of the set of SP payoffs may then be established by first constructing an SP strategy profile which generates payoff (m_i^*, n_{-i}^*) , and then demonstrating that any feasible outcome may be supported by SP strategies if and only if its expected payoff at time 0 lies between m^* and n^* .

Theorem 7: An outcome (η, τ) is SP if and only if $n^* \leq E[\beta^r \eta | \sigma_0 = s] \leq m^*$. Moreover, for each $i \in K$, there is a stationary SP outcome, (μ^i, S^{μ^i}) such that $v_i^{\mu^i} = m_i^*$ and $v_{-i}^{\mu^i} = n_{-i}^*$.

We illustrate the type of strategy required to support an arbitrary SP payoff in the discussion following Theorem 8 below.

Given Theorem 7, the problem of determining the conditions for the uniqueness of an SP payoff may be reduced to the problem of determining the conditions under which $m^* = n^*$. For a stationary nonstochastic game with k players and transferable utility, Herrero (1985) demonstrates that the SP outcome is unique when the discount factor is less than $1/(k-1)$. Merlo and Wilson (1992) extend this result to a larger class of stochastic bargaining models which include the games studied in this paper as a special case.

Theorem 8: Suppose $\beta < 1/(k-1)$. Then there is a unique SP payoff.

Note that Theorem 8 implies that the SP payoff is unique for any two-player game. For k -player games, however, the upper bound on β cannot be relaxed in the absence of any further restrictions. Sutton (1986) reports the following observation due to A. Shaked. Given a 3-player game with cake size 1 in each period in which the order of moves cycles from period to period, any feasible allocation may be sustained as an SSP outcome whenever the players have a fixed identical discount factor $\delta \geq 1/2$. Since the proof of Theorem 8 is based on a generalization of Shaked's argument, it is instructive to describe the strategy profile which supports the extremal SP payoffs.

The idea is to construct a strategy profile which employs the allocations $(1,0,0)$, $(0,1,0)$, and $(0,0,1)$ as mutually credible threats. In each period, the players employ one of three local strategies, determined by the current "state" $i \in \{1,2,3\}$ of the history of actions.⁴ In state i , the SP strategy profile requires the proposer to offer player i (possibly himself) all of the cake and for each of the responders to accept the proposal. If the prescribed allocation is proposed, but is rejected by player $j \neq i$, then the game remains in state i in the next period. If the proposer offers some other allocation a , then, for one of the responders j , $a_j \leq 1/2$, and the game moves to state j in the next period. It is a simple exercise to verify that the last player to deviate always earns a return of 0, and, consequently, the

⁴ We put "state" in quotes to distinguish from an element of S which determines the current proposer and cake size in a stochastic game.

prescribed strategy is subgame perfect regardless of the state in which the game begins.

We conclude this section with the observation that, since an SSP payoff always exists, a unique SP payoff is also an SSP payoff.

6. Concluding Remarks

In this paper we extended the standard sequential bargaining model with complete information and transferable utility to allow both the size of the surplus and the identity of the proposer to follow an arbitrary stochastic process. We provided a general characterization of both the subgame perfect and stationary subgame perfect payoffs for an n-player game and established the existence of a unique stationary subgame perfect payoff. We also exploited the stochastic elements of the game to analyze several issues which cannot be addressed with a deterministic model such as the conditions under which agreement is delayed, how the processes determining the proposer and the cake size affect the bargaining power of the individual players, and how the equilibrium payoffs are affected when the identity of the proposer is correlated with the size of the cake.

A critical assumption in much of our analysis is that utility is transferable so that the total payoff is always well defined. Although this assumption does not necessarily imply that players are risk neutral with respect to the underlying physical allocations or physical flows, it is still highly restrictive. In Merlo and Wilson (1992) we extend the framework presented here to allow for nontransferable utility. In that paper, we obtain similar characterizations of the subgame perfect and stationary subgame perfect payoffs. However, many of the more specific results such as the

uniqueness of stationary subgame perfect payoffs for n-player games must be qualified.

A more serious restriction of the model, as in almost all existing models of bargaining, is that the surplus to be allocated is taken as primitive. Consequently, a number of important issues which depend on the properties of the underlying contracts, such as recontracting, cannot be adequately explored. Nevertheless, once the implications for the process of cake size are determined from a more complete model of the underlying environment, our analysis does provide a method for determining the structure of the equilibrium.

References

- Abreu, D. (1988), "On the Theory of Infinitely Repeated Games with Discounting," Econometrica, 56, 383-396.
- Admati, A. R. and M. Perry (1987), "Strategic Delay in Bargaining," Review of Economic Studies, 54, 345-364.
- Baron, D. P. and J. A. Ferejohn (1989), "Bargaining in Legislatures," American Political Science Review, 83, 1181-1206.
- Baron, D. P. and E. Kalai (1990), "Dividing a Cake by Majority: the Simplest Equilibria," Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion Paper # 919.
- Binmore, K. G. (1987), "Perfect Equilibria in Bargaining Models," pp. 77-106, in K. G. Binmore and P. Dasgupta (eds.), The Economics of Bargaining, Oxford: Basil Blackwell.
- Binmore, K. G., A. Rubinstein and A. Wolinsky (1986), "The Nash Bargaining Solution in Economic Modelling," Rand Journal of Economics, 17, 176-188.
- Binmore, K. G., A. Shaked and J. Sutton (1989), "An Outside Option Experiment," Quarterly Journal of Economics, 104, 753-770.
- Blackwell, D. (1965), "Discounted Dynamic Programming", Annals of Mathematical Statistics, 36, 226-235.
- Chae, S. and J. Yang (1988), "The Unique Perfect Equilibrium of an N-Person Bargaining Game," Economic Letters, 28, 221-223.
- Chatterjee, K. and L. Samuelson (1990), "Perfect Equilibria in Simultaneous Offer Bargaining," International Journal of Game Theory, 19, 237-267.
- Haller, H. (1986), "Non-Cooperative Bargaining of $N \geq 3$ Players," Economic Letters, 22, 11-13.
- Herrero, M. J. (1985), "A Strategic Bargaining Approach to Market Institutions," Ph.D. Dissertation, London School of Economics.
- Karlin, S. and H. M. Taylor (1975), A First Course in Stochastic Processes, New York: Academic Press.
- Krishna, V. and R. Serrano (1991), "Multilateral Bargaining," unpublished manuscript, Harvard University.
- Merlo, A. (1992), "Bargaining over Governments," Ph.D. Dissertation, New York University.
- Merlo, A. and C. Wilson (1992), "A Stochastic Model of Sequential Bargaining with Complete Information and Nontransferrable Utility," New York

- University, C.V. Starr Center for Applied Economics, Research Report # 92-35.
- Moulin, H. (1986), Game Theory for the Social Sciences, New York: New York University Press.
- Osborne, M. J. and A. Rubinstein (1990), Bargaining and Markets, San Diego: Academic Press.
- Perry, M. and P. J. Reny (1989), "A Non-Cooperative Bargaining Model with Strategically Timed Offers," Hebrew University, Department of Economics, Working Paper # 218, forthcoming in The Journal of Economic Theory.
- Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," Econometrica, 50, 97-109.
- Sakovics, J. (1990), "Delay in Bargaining Games with Complete Information," Stanford University, Stanford Center on Conflict and Negotiation, Working Paper # 15, forthcoming in The Journal of Economic Theory.
- Shaked, A. and J. Sutton (1984a), "Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model," Econometrica, 52, 1351-1364.
- Shaked, A. and J. Sutton (1984b), "The Semi-Walrasian Economy," London School of Economics, ST/ICERD, Discussion Paper # 84/98.
- Sobel, J. and I. Takahashi (1983), "A Multistage Model of Bargaining," Review of Economic Studies, 50, 411-426.
- Stahl, D. O. II (1990), "Bargaining with Durable Offers and Endogenous Timing," Games and Economic Behavior, 2, 173-187.
- Stahl, I. (1972), Bargaining Theory, Stockholm: Economic Research Institute, Stockholm School of Economics.
- Sutton, J. (1986), "Non-Cooperative Bargaining Theory: An Introduction," Review of Economic Studies, 53, 709-724.