Markov Perfect Equilibrium in a Stochastic Bargaining Model

Branislav L. Slantchev*
Department of Political Science
University of California, San Diego

November 30, 2002

Abstract

I present a model in which two players bargain using the alternating-offers protocol while costly fighting goes on according to a stochastic process that moves some player closer to complete victory. There are many Nash equilibria and a large range of payoffs can be supported in equilibrium. However, there is a unique Markov perfect equilibrium, which is efficient, and in which offers depend on players' prospects in war as well as the current military position.

1 Introduction

In 1968, reviewing ten years of contributions to the *Journal of Conflict Resolution*, Elizabeth Converse remarked, "I get the feeling that, for most JCR contributors, once a war happens, it ceases to be interesting" (Converse 1968, pp. 476-7). Recently, political scientists have begun increasingly to address this vexing problem.[1]

Before we can begin exploring the incentives for ending wars, a model of wartime negotiations is needed. I develop a basic model that is intended to serve as a baseline for exploring the various issues that impinge on the warring parties' incentives to resolve the conflict. I analyze the model under complete information to demonstrate that the stochastic element of warfare, that is, the probabilistic nature of outcomes on the battlefield victories, is not sufficient to explain why wars continue. I show that, contrary to common arguments that blame the expansion of demands following victories for the inability to reach a settlement, this expansion is always balanced and outweighed by the incentives to conclude the costly conflict as soon as possible. I also find that the shadow of the future has somewhat surprising effects on the behavior of states that are nearing defeat.

I model war as a bargaining process, where parties alternate making proposals and counter-proposals, and where fighting occurs while disagreement persists. However, this process cannot continue indefinitely because eventually one side will win the war, although the timing and the winner are determined probabilistically. As war progresses, both sides observe how well they are doing so far and evaluate their prospects of the future. Under complete information, I find that there is a unique Markov perfect equilibrium, in which states immediately settle. Thus, although bargaining is superimposed on a stochastic warfare process, if states agree in their expectations, then they can reach a mutually satisfactory agreement.

2 Related Literature

The strategic approach to bargaining theory, initiated by Ståhl (1972) and Rubinstein (1982), is particularly suited to modeling situations where players' reversion points (payoffs in case of breakdown of bargaining) and the size of the benefits to be distributed can evolve endogenously and/or stochastically. These models generally have efficient subgame perfect equilibria under complete information, in which agreement is reached immediately.[2]

The literature on strategic bargaining in stochastic environments is relatively new. Merlo & Wilson (1995) provide a general n-player infinite-horizon complete information bargaining model, in which the identity of the proposer and the size of the pie follow a general Markov process. They find some that some subgame perfect equilibria are inefficient because players may delay agreement in the expectation that the benefits to be divided will increase.

Furasawa & Wen (2001) consider a model in which the interim disagreement payoff is determined stochastically in each period, and where the proposer can choose to delay making an offer. They find a unique perfect equilibrium that is inefficient because of a stochastically delayed agreement. The ability to delay agreement is a necessary condition for existence of the inefficient equilibrium.[3]

There is a small, but expanding, formal literature on bargaining while fighting. Although he did not analyze the model he proposed, Wagner (2000) was the first to use a Rubinstein-type alternating-offers divide-the-pie game with an exogenous risk of breakdown to model the dynamics of intrawar bargaining (Binmore, Rubinstein & Wolinsky 1986, Osborne & Rubinstein 1990).

Powell (2001) proposes an alternative specification, where after an offer is rejected, states can decide whether to fight or not. If either state attacks, then the game ends with some exogenous probability; otherwise, states pay costs for fighting. Essentially, this approach transforms the war outcome into an endogenous inside option (Muthoo 1999). He distinguishes between a "dissatisfied" state, which prefers fighting to the finish to living with

*Email: slantchev@ucsd.edu.

[1] There is some work done in a less theoretical vein using various perspectives including rationalist, psychological, bureaucratic, and Marxist, by Carroll (1969), Fox (1970), Halperin (1970), Kecskemeti (1970), Foster & Brewer (1976), Handel (1978), Oren (1982), Manwaring (1987), and Engelbrecht (1992). Pillar (1983) provides a useful bargaining view of war termination, and Goemans (2000) presents a thoughtful study of the politics of war termination during the First World War.

[2] See Muthoo (1999) for a an excellent review.

[3] The model in this paper can be conceptualized as one with variable interim disagreement payoffs (e.g. if we interpret the expected payoff from fighting in Definition 4.1 in this way), and thus it is not surprising that in our case a unique efficient perfect equilibrium always exists, confirming the necessity condition described above.

the status quo, and a "satisfied" one, which does not. The complete information equilibria are efficient and the optimal offer depends on how dissatisfied a state is. Powell then analyzes two possible forms of uncertainty: about the distribution of power, and about the opponent's costs of fighting. This permits a characterization of the screening process on the basis of the time it takes to marshal resources in order to fight, as opposed to signaling with an offer. Generally, he finds that the results of previous costly-lottery models carry over to war as a process models.

Filson & Werner (2002) model the process with alternating phases of negotiations and fighting in case no agreement is reached. Initially, states control certain amounts of resources and consume benefits. Resources enter the payoff functions directly but are also important for a state's ability to continue fighting because if resources fall below some threshold level, then a state is assumed defeated and loses its benefits. There is one-sided uncertainty about the probability of winning one battle, which results in a screening equilibrium, in which outcomes depend on the attacker's beliefs about the type of defender.

The model presented here is an extension of the stochastic model of warfare analyzed by Smith (1998), where the two states fight over "forts" and derive utility from the number of forts in their possession. Because controlling more forts increases the payoffs directly, the model in effect makes war profitable by assumption.

Kim (2001) builds upon this work by modeling the two states as bargaining over the division of a finite number of districts of military and intrinsic value.

There is asymmetric information about player 2's fighting costs, and only the uninformed player is allowed to make offers. In the unique perfect Bayesian equilibrium, the familiar screening process occurs, in which high cost types of player 2 drop out sooner by accepting smaller offers. The model presented here is related to this work as well, although it admits discounting, offers a completely different way of conceptualizing uncertainty, assumes no intrinsic value of the object contested militarily, and does not restrict the role of the proposer.

Smith & Stam (2001) provide an modification of this model that presents the most appealing formulation that is designed to account for the instrumental role of warfare. They superimpose a one-sided bargaining process according to which states negotiate a division of benefits while fighting. Players are infinitely patient (there is no discounting) and only one player is allowed to make offers. Somewhat controversially, the authors introduce asymmetric information in a non-standard way. They posit that players have divergent beliefs over the eventual outcome of the war and then analyze the model by ignoring the crucial information conveyed by the behavior of the opponent.

I also analyze warfare as an instrument used in pursuit of political (distributional) objectives, and so my model derives from, and extends, that of Smith & Stam (2001). I introduce two important modifications that turn out to be crucial: states have time preferences (they discount the future), and the bargaining protocol allows both sides to make offers. In a related paper, I analyze this model under asymmetric information using standard game-theoretic techniques and arrive at conclusions strikingly different from those of Smith & Stam (2001). Both under complete information and under uncertainty, the extensions of their model drive the results to a large extent.

3 The Model

Two players, $i \in \{1, 2\}$, bargain over a two-way partition of a flow of benefits with size π. An agreement is a pair (x, y), where x is player 1's share, and y is player 2's share. The set of possible pairs is

$$\mathcal{X} = \left\{ (x, y) \in \mathcal{R}^2 : x + y = \pi \text{ and } 0 \le x, y \le \pi \right\}.$$

Players have strictly opposed preferences and each is concerned only with the share of benefits it obtains from the agreement. Because a share x identifies a distribution uniquely, let x be equivalent to the pair $(x, \pi - x)$, and y be equivalent to the pair $(\pi - y, y)$. The status quo distribution of benefits is (s_1, s_2) with $s_1 + s_2 = \pi$.

The two players bargain according to the alternating-offers protocol (Rubinstein 1982). Players have a common discount factor $\delta \in (0, 1)$, and act in discrete time with a potentially infinite horizon and periods indexed by t ($t = 0, 1, 2, \ldots$). In even-numbered periods, player 1 proposes a division $x \in \mathcal{X}$ to player 2. If player 2 accepts that proposal, an agreement is reached, and the game ends with players receiving their shares in $(x, \pi - x)$. If player 2 rejects the proposal, then players fight a costly engagement, which may improve the relative military position of a player, and the period ends. Player 2 makes a counteroffer $y \in \mathcal{X}$ in the next period. If player 1 accepts, the game ends and players receive their payoffs from the agreement $(\pi - y, y)$; if it rejects, they fight another military engagement. The game continues until an agreement is struck or until one of the players is decisively defeated. If a player decisively defeats the other, then it obtains the entire flow of benefits π. Each military engagement is costly, and states suffer a constant per-period loss of utility, reducing their instantaneous per-period wartime payoffs to $b_i < s_i$ (and so $b_1 + b_2 < \pi$).

War is modeled as a stochastic process of attrition. It is a homogenous Markov chain with two absorbing states: victory and loss.[4] The *current military position* of a player at time t captures the player's relative overall success from all engagements that have occurred up to time t. Let $N \ge 2$ denote the finite number of military objectives and let k be the number of objectives achieved by player 1. The set of possible states is $K = \{0, 1, \ldots, N\}$. At time t, player 1's current military position, $k_t \in K$, is the difference between the total number of its victories and losses in battles that have occurred in periods $(0, 1, \ldots, t - 1)$. The state variable k_t is an indicator of relative military advantage at time t and summarizes the whole history of the war up to that point in time; k_0 is the position at the outset of war.

One battle over one objective occurs in each period. Player 1 wins the fight with probability p, and loses with probability $1 - p$.[5] If player 1 wins the battle at time t, then $k_{t+1} = k_t + 1$, and if it loses, then $k_{t+1} = k_t - 1$. If $k_t = 0$, player 1 is militarily defeated and the game ends with player 2 imposing the settlement $(0, \pi)$. If $k_t = N$,

[4] See Grimmett & Stirzaker (1992) and Norris (1997) for random processes.

[5] It is possible to relax this assumption in two ways: the probabilities of victory and defeat need not sum to 1; the probabilities of victory and defeat vary according to the current military position. It is not clear *a priori* how probabilities should vary with battlefield success. For example, sometimes early failure mobilizes the will to fight (e.g. Britain after Dunkirk in May 1940) but other times it does not (e.g France in June 1940).

player 2 is militarily defeated and player 1 imposes the settlement $(\pi, 0)$.

Let $\mathbf{P} = (q_{ij} : i, j \in [0, N])$ be the stochastic matrix, where $q_{ij} = \Pr(k_{t+1} = j : k_t = i)$. Clearly, the transition matrix $\mathbf{P}$ is square $(N+1) \times (N+1)$. The probabilities of victory and defeat described above induce a probability distribution for each row of $\mathbf{P}$ such that

$$q_{ij} = \begin{cases} p & \text{if } j = i+1 \text{ and } 0 < i < N \\ 1-p & \text{if } j = i-1 \text{ and } 0 < i < N \\ 1 & \text{if } i = j \text{ and } i \in \{0, N\} \\ 0 & \text{otherwise} \end{cases}$$

Let $\theta(\mathbf{P})$ denote the temporally homogenous Markov process with stochastic matrix $\mathbf{P}$ realizing values in $K = \{0, 1, \ldots, N\}$. For $t = 0, 1, \ldots$, let $\theta^t \equiv (\theta_0, \theta_1, \ldots, \theta_t)$ denote the state history at time t with typical realization $(k_0, k_1, \ldots, k_t)$ where the elements $k \in K$ are referred to as *states*.

The stochastic bargaining game is played as follows. Player 1 is the proposer in all even periods and responder in all odd periods, when player 2 is the proposer. In each period t, a state k is realized and the proposer i offers a partition in $\mathcal{X}$. The other player responds by either accepting or rejecting the proposal. If the offer is rejected, both players obtain their per-period instantaneous payoffs b_i, and a new state is realized in the next period according to $\theta(\mathbf{P})$. The stage game continues until either a partition is accepted or one of the absorbing states is realized.

A strategy specifies the offer that a player must make when it is the proposer, and its reaction to any offer that its opponent makes. In every period t, the information set for the proposer consists of the history of rejected offers and realizations of the state variable. A pure behavioral strategy specifies the current proposal as a function of this history. In the same period, a pure behavioral strategy for the responder specifies a function from this history concatenated with the current offer to an action in the set $\{Y, N\}$, which prescribes either acceptance or rejection of the current offer.

To specify the strategies formally, let h_t denote the history of play up to period t. Recall that if t is even, then player 1 is the proposer. In what follows, I describe the strategies for player 1, the strategies for player 2 can be specified analogously. Let x_t denote the partition $(x, \pi - x)$ that player 1 offers at time t, and let $y_{t'}$ denote the partition $(\pi - y_{t'}, y_{t'})$ that player 2 offers at time t'.

At some even date t player 1 knows $h_t = \left\{\langle k_0, x_0\rangle, \langle k_1, y_1\rangle, \ldots, \langle k_{t-1}, y_{t-1}\rangle, k_t\right\}$, which consists of t realizations of the state variable (not counting the initial state k_0), and $t-1$ previous offers and counteroffers. At some odd date t', the history $h_{t'}$ consists of $h_{t'-1}$ concatenated with $(x_{t'-1}, \langle k_{t'}, y_{t'}\rangle)$, that is, player 1's offer in the previous period, the new realization of the state variable, and the current offer of player 2. Since the game only continues while players reject proposals or until one of the absorbing states is realized, rejections are omitted from the specification. Taking H_t to be the set of all such histories, a pure behavioral strategy for player 1, Σ^1, is a sequence of functions $\{\sigma_t^1\}_{t=0}^{\infty}$ such that $\sigma_t^1 : H_t \to \mathcal{X}$ when t is even, and $\sigma_t^1 : H_t \to \{Y, N\}$ when t is odd. For example, $\sigma_2^1(\langle k_0, x_0\rangle, \langle k_1, y_1\rangle, k_2)$ specifies player 1's offer at time $t = 2$ assuming that it offered x_0 in period 0, its opponent rejected it and countered with y_1, which player 1 rejected, and the realizations of the state variable were k_1, and k_2. Note that I only consider pure strategies throughout this article.

If players adopt strategies σ^1 and σ^2, the outcome of the game is $R(\sigma^1, \sigma^2) = \langle R^1(x, t), R^2(x, t)\rangle$ where $R^i(x, t)$ specifies player i's expected payoff if agreement on partition x is reached at time t, and where (x, t) are determined from the strategies. For example, suppose the strategies specify that players reject all offers until time T, at which point some x is proposed and accepted. The payoffs then will be $R^i(x, T)$, which, loosely speaking, represent the payoff from partition x agreed on in period T times the probability that the game continues until T plus the expected payoff if the game ends prior to T times the probability that it ends. The precise definition of $R^i(\cdot)$ is somewhat convoluted and is relegated to Appendix B.[6]

Players maximize the time-averaged discounted sum of per-period payoffs:

$$(1-\delta) \sum_{t=0}^{\infty} \delta^t r_t^i$$

where r_t^i is player i's instantaneous payoff in period t, which equals b_i if players disagree, 0 if player i loses the war, π if it wins, and i's share of benefits if players terminate the war with a settlement.

It is worthwhile to note that this model avoids some pitfalls associated with prevalent ways of modeling war. First, unlike the costly-lottery approach, it does not reduce war to a single-shot event, and permits analysis of dynamics. Second, unlike the infinitely repeated game approach, as commonly used, it does not go against the intuition that the process does not last indefinitely, or even a large number of periods.[7] Instead, this model captures the dynamic nature of the process without either fixing an arbitrary number of periods or allowing it to extend indefinitely, while incorporating the time-dependence of each state.

4 Nash Equilibria

Perpetual disagreement will be denoted by $(\cdot, \infty)$. In this case players receive their per-period disagreement payoffs while fighting lasts and then the winner receives the entire flow of benefits forever after the war ends.

Definition 4.1. The function $W_k^i : K \to [0, \pi]$ is player i's payoff from fighting to the finish starting in state k. It is defined recursively for $0 < k < N$:

$$W_0^1 = W_N^2 = 0$$
$$W_N^1 = W_0^2 = \pi$$
$$W_k^i = (1-\delta) b_i + \delta\left[p W_{k+1}^i + (1-p) W_{k-1}^i\right]$$

Appendix A shows that W_k^i has a closed form. Since $0 \le b_i < \pi$, it follows that $W_k^i \in [0, \pi]$ for all k. Provided that the strategies are such that they specify perpetual disagreement, the outcome is $\langle R^1(\cdot, \infty), R^2(\cdot, \infty)\rangle = \langle W_{k_0}^1, W_{k_0}^2\rangle$.

[6] In addition to the derivations and proofs in the appendices, several computer programs (in C++and Gauss) that perform numerical simulations and compute the examples are available from the author upon request.

[7] For an especially illuminating discussion of this topic, see Rubinstein (1991, p. 918).

The following claim, which is straightforward to verify, establishes an important property of the stochastic bargaining game.

CLAIM 4.1. *$W_k^1 + W_k^2 < \pi$ for any $\delta \in (0,1)$ and all $k \in K$.*

That is, there is no state in which players lack incentives to bargain. Since the payoffs players expect to get if war is fought until the end sum to less than the total benefits they can redistribute, there always exists a surplus that can be allocated to make both players better off with a bargain. That is, there are always gains from a negotiated agreement.

The set of Nash equilibria for the bargaining game with stochastic fighting under complete information is very large. Instead of attempting to characterize all Nash equilibria, I shall only provide a characterization of the set of equilibrium payoffs. Perhaps not surprisingly, every distribution of benefits that is better for both players than their expected payoffs from fighting to the bitter end from the initial state can be supported by some Nash equilibrium.

LEMMA 4.1. *There exists no equilibrium, in which player i's payoff is less than $W_{k_0}^i$.*

Proof. Suppose some equilibrium yielded the expected payoff $x < W_{k_0}^i$ to player i. The only equilibria that may produce such an outcome are those where the game ends with an agreement at some time $t < \infty$. However, player i can deviate to a strategy that rejects all offers and demands the entire π every time i gets to make a proposal. The expected payoff from this strategy is $W_{k_0}^i$, making a deviation profitable. This contradicts the supposed optimality of the original strategy, and establishes the claim. Q.E.D.

Lemma 4.1 establishes the bounds on the payoffs that can be supported in equilibrium. Expressed in terms of player 1's share, the range of equilibrium payoffs is thus $x \in [W_{k_0}^1, \pi - W_{k_0}^2]$.

PROPOSITION 4.2. *Any partition x is a Nash equilibrium outcome if and only if*

$$x \in [W_{k_0}^1, \pi - W_{k_0}^2]$$

Proof. Necessity follows from Lemma 4.1. To establish sufficiency, we must show the existence of strategies that yield x in equilibrium provided x is in the prescribed range.

Consider the following simple strategies: both players demand π and reject all offers except in period t, when player 1 demands $\hat{x}$ and player 2 accepts it. These simple strategies are very attractive because there is only one opportunity to strike a bargain, and so there is only one interesting deviation to examine, all others being trivial. This is the deviation where player 1 does not offer (or player 2 does not accept) x. Because such a deviation yields the expected payoff from fighting to the bitter end, comparisons are easy.

We must find a pair $\langle \hat{x}, t \rangle$ such that $R^1(\hat{x}, t) = x$ and $R^2(\hat{x}, t) = \pi - x$. It is readily verified that the pair $\langle x, 0 \rangle$ satisfies both conditions, and so the strategies with payoffs $\langle R_1(x,0), R_2(x,0) \rangle$ constitute a Nash equilibrium. Q.E.D.

Proposition 4.2 implies that every distribution that falls within the range specified by Lemma 4.1 can be supported in a Nash equilibrium where player 1 proposes it only in the initial period and then always demands π and rejects all offers. Given such a strategy, player 2 can do no better than accept $\pi - x$ immediately. Conversely, given that player 2 always demands π and rejects all offers except $\pi - x$ in the initial period, player 1 can do no better.

Such Nash equilibria have the unattractive property that they are not subgame perfect. In particular, they require that players threaten to reject any offer, even ones that will be fairly attractive if made at opportune moments. For example, suppose player 2 deviates and rejects 1's initial offer. Following the Nash strategies requires now fighting to the bitter end. However, suppose war develops in such a way that player 1 comes very close to defeat, i.e. $k = 1 < k_0$. Can it credibly threaten to reject any offer? Suppose player 2 offers it a little more than the certainty equivalent from fighting from the new state. Accepting such an offer is strictly better for player 1 from its current vantage point. Thus, the threat to reject it is not credible. Nash equilibrium, however, has no bite for such considerations because it only requires optimality on the equilibrium path. In the following section, I examine the existence of an equilibrium where threats and promises are credible in the sense that if players must implement them, it will be in their interest to do so. The subgame perfection refinement requires that play be optimal whether or not it occurs on the equilibrium path.

5 The Markov Perfect Equilibrium

Given the model of warfare described the previous section, it seems natural to restrict attention to bargaining strategies where offers and acceptance rules depend on the current military position and the expectations about the eventual outcome and duration of war. These strategies, which condition on the realization of a *state variable* are commonly known as Markov. In this model, the only variable that influences current and future payoffs is the current military position. Because the class of Markov strategies requires that equilibrium depends only on payoff-relevant history, players' strategies at time t must depend only on k_t.[8]

I shall be looking for Markov perfect equilibria, that is, perfect equilibria in Markov strategies (Fudenberg & Tirole 1991). A pair of strategies forms a *Markov perfect equilibrium* (MPE) if, and only if, each player's strategy maximizes its intertemporal payoff at any time t, given k_t and assuming that henceforth both states play according to their respective reaction functions. In other words, a MPE is a profile of Markov strategies that yields a Nash equilibrium in every proper subgame.

This class of equilibria has several remarkable properties. First, MPE often succeeds in eliminating or reducing the multiplicity of equilibria in dynamic games. Second, it

[8] Since Nash equilibrium may rely on non-credible threats and promises off the equilibrium path, the solution concept I shall use is that of subgame perfect equilibrium, which requires that each player's strategy is optimal in every proper subgame, whether or not this subgame is ever reached when players follow their strategies (Selten 1975). In other words, an equilibrium is subgame perfect if the strategies form a Nash equilibrium in every subgame.

eliminates the problem of states reacting to military positions they had several periods in the past. Although the restriction to payoff-relevant history makes complicated punishment strategies impossible, these strategies are often hard to justify intuitively.[9]

I now derive a useful property that all MPE possess. Suppose that player 2 has to make an offer at some time t when the state is k. Since player 1 expects to get W_k^1 if it rejects the proposal and continues fighting, player 2 must offer it at least W_k^1. That is, player 2 must offer 1 at least the value of its expected resolution of war, otherwise player 1 would reject the proposal because continuing yields a larger expected payoff. However, since W_k^2 is the present value of the expected resolution of war for player 2, it follows that player 2 will never propose for itself less than this amount either. This gives rise to the following lemma, which establishes a necessary condition for the equilibrium offers.

LEMMA 5.1. *Let k be some realization of the Markov process at time t, and let x_k^* and y_k^* be the optimal offers of player 1 and player 2, respectively. Then, in any MPE,*

$$x_k^* \in \left[W_k^1, \pi - W_k^2\right] \text{ and } y_k^* \in \left[W_k^2, \pi - W_k^1\right].$$

In equilibrium neither state would accept less than its certainty equivalent for fighting, which places the bounds on equilibrium proposals. These bounds restrict offers to be decidedly different (whenever state payoffs are strictly greater than 0) from the Rubinstein model whenever a country is nearing victory or defeat. That is, while in the original alternating-offers model players split the benefits evenly when the time between offers goes to zero (or in any event the advantage of the player who moves first is captured by the discount factors), in this model military position affects expectations directly and a player close to victory will not agree to less than the expected value of victory, which may well be close to taking the entire flow of benefits. These considerations are elaborated in the text that follows.

The interval in the lemma represents the gains from concluding a negotiated settlement. From Claim 4.1 this interval always exists. Thus, not surprisingly, the following section establishes the existence of an efficient equilibrium, in which players split this surplus depending on the initial state of the world.

5.1 Existence and Uniqueness of MPE

In a *no-delay* MPE, a player's equilibrium proposal is immediately accepted by the other player. In a *stationary* MPE, players always make the same state-dependent offers (that is, offers depend on the realization of the stochastic process but are otherwise time-invariant). In any stationary no-delay MPE, player 1 must offer player 2 at least what that player expects to obtain by rejecting a proposal. Since in this equilibrium player 2's offer is immediately accepted (or else the game ends), player 1's offer must satisfy, for $1 < k < N - 1$

$$\begin{aligned} \pi - x_1^* &= (1-\delta)b_2 + \delta[py_2^* + (1-p)\pi] \\ \pi - x_k^* &= (1-\delta)b_2 + \delta[py_{k+1}^* + (1-p)y_{k-1}^*] \\ \pi - x_{N-1}^* &= (1-\delta)b_2 + \delta(1-p)y_{N-2}^* \end{aligned} \quad (1)$$

Similarly, player 2's offer must satisfy

$$\begin{aligned} \pi - y_1^* &= (1-\delta)b_1 + \delta p x_2^* \\ \pi - y_k^* &= (1-\delta)b_1 + \delta[px_{k+1}^* + (1-p)x_{k-1}^*] \\ \pi - y_{N-1}^* &= (1-\delta)b_1 + \delta[p\pi + (1-p)x_{N-2}^*] \end{aligned}$$

This defines a system of $2(N-1)$ simultaneous equations. For example, for $N = 4$, we have (letting $z_k \equiv x_k^*$ and $z_{k+3} \equiv y_k^*$ for $k = 1, 2, 3$)

$$\begin{aligned} \pi - z_1 &= (1-\delta)b_2 + \delta[pz_5 + (1-p)\pi] \\ \pi - z_2 &= (1-\delta)b_2 + \delta[pz_6 + (1-p)z_4] \\ \pi - z_3 &= (1-\delta)b_2 + \delta(1-p)z_5 \\ \pi - z_4 &= (1-\delta)b_1 + \delta p z_2 \\ \pi - z_5 &= (1-\delta)b_1 + \delta[pz_3 + (1-p)z_1] \\ \pi - z_6 &= (1-\delta)b_1 + \delta[p\pi + (1-p)z_2] \end{aligned}$$

Letting $\mathbf{z} = (z_1\ z_2\ z_3\ z_4\ z_5\ z_6)^T$, this can be compactly written in the familiar matrix form $\mathbf{Az} = \mathbf{w}$, where the coefficient matrix $\mathbf{A}$ is given by:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & \delta p & 0 \\ 0 & 1 & 0 & \delta(1-p) & 0 & \delta p \\ 0 & 0 & 1 & 0 & \delta(1-p) & 0 \\ 0 & \delta p & 0 & 1 & 0 & 0 \\ \delta(1-p) & 0 & \delta p & 0 & 1 & 0 \\ 0 & \delta(1-p) & 0 & 0 & 0 & 1 \end{pmatrix}$$

and

$$\mathbf{w} = \begin{pmatrix} \pi - \delta(1-p)\pi - (1-\delta)b_2 \\ \pi - (1-\delta)b_2 \\ \pi - (1-\delta)b_2 \\ \pi - (1-\delta)b_1 \\ \pi - (1-\delta)b_1 \\ \pi - \delta p\pi - (1-\delta)b_1 \end{pmatrix}$$

It is trivial to verify that the coefficient matrix is full rank, which implies that it is nonsingular and therefore the equation has a unique nonzero solution. Computing the closed form of each z_i expressed in terms of the exogenous parameters is possible, but is extremely tedious, and, because the resulting expressions are very long, prone to error. The complications become worse very quickly as N increases, making the direct approach impractical. Instead, I will demonstrate that for any $N \geq 2$, the equivalent system of equations always has a unique solution. I will then provide several examples with interesting values of the exogenous parameters to give some intuition of how the equilibrium offers behave.

Fix some arbitrary $N \geq 2$ and let $n = N - 1$. Label the equilibrium offers such that $x_1^*, x_2^*, \ldots, x_n^*$ correspond to $z_1, z_2, \ldots, z_n$; and $y_1^*, y_2^*, \ldots, y_n^*$ correspond to $z_{n+1}, z_{n+2}, \ldots, z_{n+n}$. Construct the $2n \times 2n$ matrix $\mathbf{A}$ of coefficients in the usual way and let $\mathbf{w} > \mathbf{0}$ be the corresponding RHS vector. The following lemma, whose proof is in Appendix C, establishes the claim.

[9] An excellent discussion of the subtle differences between Markov strategies and traditional supergame punishment strategies appears in Tirole (1993, pp.253-6). A systematic framework for analyzing infinitely repeated games with discounting appears in Abreu (1988). For an illuminating discussion of the advantages of MPE, see Maskin & Tirole (2001). It is worth noting that the Rubinstein-type bargaining models, such as this one, are not repeated games, and the traditional Folk Theorem results do not apply because deviations end the game. For a model that can support punishment during play, see **?**, where complicated strategies are made possible by the existence of a conflict game outside the bargaining process.

LEMMA 5.2. *There exists a unique* $\mathbf{z}^* = \mathbf{A}^{-1}\mathbf{w}$.

Thus, for any $N \geq 2$, the system of simultaneous equations has a unique solution, $\mathbf{z}^*$, which defines the proposals that satisfy the no-delay and stationarity properties of MPE. It now remains to be shown that there exists a pair of strategies that supports these proposals in equilibrium. The following proposition, whose proof is in Appendix C, characterizes this unique pair of Markov strategies.

PROPOSITION 5.3. *The stochastic bargaining game with complete information has a unique no-delay Markov subgame perfect equilibrium, in which player* 1 *always offers* x_k^*, *accepts all offers* $x \geq x_k^*$, *and rejects all offers* $x < x_k^*$, *where* k *is the realization of* $\theta(\mathbf{P})$, *and* x_k^* *is the* k*th element of* $\mathbf{z}^*$. *Player 2's strategy is defined analogously. Agreement is reached immediately on* $x_{k_0}^*$.

Proposition 5.3 establishes the existence of a unique stationary no-delay Markov perfect equilibrium of the bargaining game under stochastic warfare. It is clear from the construction of the optimal offers that the necessity condition of Lemma 5.1 is satisfied. Because of the alternating-offers protocol, the proposer enjoys some advantage but is not able to extract the entire surplus. This result is intuitive since the offers are moderated by the possibility that the other player may become the proposer following a rejection and this position confers an advantage. Thus, players split the surplus available in state k in a way that depends on their patience and the probability of winning.

5.2 COMPARATIVE STATICS

Since computing the closed form solution for an arbitrary N analytically is demanding and the resulting expressions unwieldy, I analyze the effect of varying the parameters of the models with computer simulations, which have the added benefit of examining the effects of simultaneous changes on the optimal offers. Even under complete information we obtain several insights into the dynamics of wartime negotiation and the incentives to conclude a bargain and end the war before the ultimate destruction of either side. I begin with a simple illustrative case.

EXAMPLE 5.1. Let $N = 10$, $\pi = 1, b_1 = .3, b_2 = .4$, with $\delta = .95$, and $p = .4$. Table 1 shows the optimal offers in each of the non-terminal states in addition to the expected payoffs from fighting to the end.

This table demonstrates three important properties of optimal offers in MPE. First, both proposal and acceptances yield strictly higher payoffs than fighting to the end. Second the player who gets to make the offer enjoys proposal power and is able to extract a larger share, as common to bargaining models that use the alternating-offers protocol. Third, payoffs are strictly increasing in state for player 1 and strictly decreasing for player 2.

For example, consider $k_0 = 6$. If player 1 gets to propose first, it offers (and player 2 accepts) $x_6^* = .359$. On the other hand, if player 2 is the proposer, player 1 will accept $\pi - y_6^* = .355$. In either case, settling is preferable to fighting to the end, which is $W_6^1 = .270$. Similar calculations hold for player 2.

An important finding is that under complete information there always exists a vector of optimal offers, which are made and accepted in the unique Markov SPE. In other words, regardless of the distribution of power, the military advantage, the status quo flow of benefits, or the costs of fighting, states always end the war with a negotiated settlement immediately.

Figure 1 shows the optimal offers depending on the military position for four cases: (a) high cost for player 1, low cost for player 2, little discounting; (b) high cost for player 1, low cost for player 2, high discounting; (c) low cost for player 1, high cost for player 2, little discounting; and (d) low cost for player 1, high cost for player 2, high discounting.

Several findings emerge from these plots. First, player 1's offers are non-decreasing in k. In other words, as its military position improves, player 1 demands (and receives) a larger share of benefits. This finding holds regardless of the probability of winning individual engagements. On the other hand, as this probability increases, so do the proposals. That is, if player 1 is more likely to win any individual battle, then, *ceteris paribus*, it will demand a larger share of benefits.

This expansion of war aims can be dramatic. Consider, for example, case (a) where, despite the higher cost of fighting, player 1's demand increases from less than .1 to .5—even when the player is close to losing the war ($k = 1$)—when p increases from .01 to .99. When player 1 is close to winning at $k = 19$, the equilibrium proposal jumps from a little over .3 when $p = .01$, to almost the entire flow of benefits at $p = .99$.

Perhaps not surprisingly, when the balance of power does not favor either player (p is close to .5), then there is a range of military positions (roughly $6 \leq k \leq 12$), where player 1 does not expand its demands much. This is evident in the nearly flat spot on the surface of all graphs. The reason for this is intuitive: With the balance of power at parity, gains on the battlefield do not translate into corresponding bargaining strength as long as the military position is not too advantageous for the player.

The second finding concerns the effect of fighting costs, or, more generally, the costs of disagreement. Consider cases (a) and (c) in Figure 1. In the first example, player 1's cost is relatively high: no benefits while fighting lasts, while player 2 is able to obtain some nonzero utility while fighting. Even though player 1 is close to complete defeat at $k = 1$, it can demand a strictly positive share of benefits. That is, the eventual settlement yields a flow of benefits that is larger than the wartime payoff. This holds even if the player is unlikely to prevail in individual battles. On the other hand, in the second example, player 1's costs are relatively low: it receives a wartime payoff of .3, while its opponent receives nothing. However, when player 1 is close to defeat *and* is not very likely to win individual engagements, it will settle for *less* than its wartime payoff. This situation is the reverse of the one observed in the previous case. Here, it is player 2 who is nearing complete victory and, since it values the future highly, can afford to absorb the high costs of fighting and therefore demand larger concessions from player 1. Because defeat is close and likely, and because it also values the future highly, player 1 is induced to concede. However, player 1's bargaining position is still stronger in case (c) than in case (a), which can be clearly seen for any position k and any distribution of power p.

The third finding is the quite significant effect of discounting. Although the present model does not allow us to assess the impact of different discount factors, we can nevertheless draw conclusions about the effect of valuing

Table 1: MPE Offers and Payoffs to Fighting to the End.

k	x_k^*	$\pi - y_k^*$	W_k^1	y_k^*	$\pi - x_k^*$	W_k^2
1	.083	.069	.052	.931	.917	.900
2	.142	.140	.096	.860	.858	.815
3	.204	.191	.137	.809	.796	.742
4	.251	.247	.176	.753	.749	.677
5	.306	.295	.218	.706	.694	.617
6	.359	.355	.270	.645	.641	.556
7	.435	.424	.344	.576	.565	.484
8	.539	.535	.461	.465	.461	.387
9	.715	.702	.658	.298	.285	.241

the future on bargains. Figure 2 illustrates the effect under varying disagreement costs, holding the distribution of power favoring player 1 significantly at $p = .8$.

Another remarkable fact is the somewhat counterintuitive role played by the discount factor in the calculations of the players. Generally speaking, if a player expects to do well in the war, either because its probability of winning a particular fight is high or because it is near to total victory, then reducing the discount factor lowers its expected payoff from fighting until the end. In other words, a strong player that is impatient would settle for less than a more patient one. Conversely, its weak opponent benefits in this case. More generally, if a player does not expect to do well in the war, either because its probability of winning individual fights is low or because it is close to defeat, being less patient improves its expected payoff from fighting until the end.

Of particular interest here are the plots for the high discount factors $\delta = .63$ and $\delta = .9$. Notice how valuation of the future works against player 1 in situations where defeat is likely but how it reverses its effect, which is especially pronounced when disagreement costs actually favor the other player. The intuition is straightforward here: When players evaluate the future almost as much as the present, then disagreement costs do not play a large role when the military resolution of war is at hand. That is, when players care much about the eventual outcome, and that outcome is favorable, the interim payoffs do not matter much because they are outweighed by the benefits of the outcome forever after. If, conversely, players care more about the present, then the proposals are mainly determined by the disagreement costs and the balance of power, and do not depend much on the military situation. This is clear from the nearly flat plots in those cases. Such situation is hard to imagine substantively because it would imply that neither player is very much concerned with the payoffs from obtaining peace. The rest of this paper, therefore, will analyze the case where players care much about the future.

It is worth noting the particularly great expansion of war aims when the probability of winning is high, when total victory is near (or, conversely, the dramatic contraction when probability of winning is low and total defeat is close), and when players discount the future significantly. With such wartime dynamics, it should be clear that the player on the losing side has incentives to seek a negotiated settlement well before the war becomes clearly not winnable. On the other hand, the winning player's demand depend on the disagreement costs, and thus this player also has incentives to settle sooner in order to avoid the costs associated with prolonging the conflict.

6 Conclusion

Lemma 5.1 established that in any MPE players never offer or accept less than what they expect to obtain if they fight the war to its bitter end. Because conflict is costly, the expected payoffs from the ultimate resolution always sum to less than the flow of benefits, which gives both players incentive to negotiate an agreement immediately (Claim 4.1). This result formalizes the intuition that fighting can cease as soon as both sides agree on the expected outcome of war. Remarkably, it closely matches a claim made almost 200 years ago:

> Not every war need be fought until one side collapses [...] Of even greater influence on the decision to make peace is the consciousness of all the effort that has already been made and of the efforts yet to come. Since war is not an act of senseless passion but is controlled by its political object, the value of this object must determine the sacrifices to be made for it in *magnitude* and also in *duration*. Once the expenditure of effort exceeds the value of the political object, the object must be renounced and peace must follow.
>
> We see then that if one side cannot completely disarm the other, *the desire for peace on either side will rise and fall with the probability of further successes and the amount of effort these would require.* If such incentives were of equal strength on both sides, the two would resolve their political disputes by meeting half way. If the incentive grows on one side, it should diminish on the other. Peace will result so long as their sum total is sufficient—though the side that feels the lesser urge for peace will naturally get the better bargain (Clausewitz 1832, pp. 91-2, emphasis mine).

Demands may or may not be constrained by the military situation: If war is sufficiently costly, then the bargaining range is virtually unconstrained. Thus, winning battles does not automatically translate into bargaining advantage. Therefore, under certain conditions, one cannot expect that victors necessarily expand their war aims; and, more importantly, unless sufficiently close to victory, states with great military advantage cannot translate it into a better bargaining position.

The complete information case can be treated as the situation in which both players agree on the probability of winning battles and are also perfectly aware of each other's costs. The efficient result stated in Proposition 5.3 confirms the intuition that wartime negotiations will end fighting whenever adversaries agree on the possible outcome. Note in particular that even though warfare follows a stochastic process, the agreement on its parameters allows both players to "predict" the likely outcome and therefore offer acceptable bargains. This makes outcomes "irreversible" in the sense that both players can calculate the expected payoffs and agree on them, giving formal support to Blainey's (1988) claim that wars end when both sides agree on their relative power.

Simply because war is a sequence of battles with probabilistic outcomes does not imply that states will have any difficulty finding mutually satisfactory bargains. Relative strength, as expressed in the probability of winning individual battles and in the costs of disagreement, determines the bargains that are mutually satisfactory (in the sense that they are preferable to fighting). Rationality then requires that states conclude such bargains even if they do not like the terms, which is what happens in the unique equilibrium. This suggests that we should look at problems arising from lack of complete information.

In models of intrawar bargaining, the central function of fighting is to reveal private information either about the type of opponent or about the distribution of power. These models generally predict that as fighting continues and states update their beliefs about the relevant parameters, settlement becomes more likely. The model with endogenous war aims, on the other hand, posits that as states draw closer to victory, they may expand their demands in reflection of their relative military advantage gained from battlefield success. In the unconstrained model here, the losing side can always settle by offering the necessary amount to make its opponent indifferent between fighting to a total victory and accepting the proposed division. On the other hand, it is not difficult to imagine how constraints on what can be offered, e.g. domestic or ideological, may prevent the losing side from making an optimal offer, in which case it may prolong fighting in the hope that its fortunes of war will reverse.

References

Abreu, Dilip. 1988. "On the Theory of Infinitely Repeated Games with Discounting." *Econometrica* 56(March):383–96.

Binmore, Kenneth G., Ariel Rubinstein & A. Wolinsky. 1986. "The Nash Bargaining Solution in Economic Modelling." *Rand Journal of Economics* 17(Summer):176–88.

Blainey, Geoffrey. 1988. *The Causes of War.* 3rd ed. New York: The Free Press.

Carroll, Berenice A. 1969. "How Wars End: An Analysis of Some Current Hypotheses." *Journal of Peace Research* 4:295–321.

Clausewitz, Carl von. 1832. *On War.* Princeton: Princeton University Press. Edited and translated by Michael Howard and Peter Paret, 1984.

Converse, Elizabeth. 1968. "The War of All Against All: A Review of the Journal of Conflict Resolution, 1957-68." *The Journal of Conflict Resolution* 12(December).

Engelbrecht, Jr., Joseph A. 1992. *War Termination: Why Does a State Decide to Stop Fighting.* New York: Columbia University. Doctoral dissertation.

Filson, Darren & Suzanne Werner. 2002. "A Bargaining Model of War and Peace: Anticipating the Onset, Duration, and Outcome of War." *American Journal of Political Science* 46(October):819–838.

Foster, James L. & Garry D. Brewer. 1976. "And the Clocks were Striking Thirteen: The Termination of War." *Policy Sciences* 7(June):225–43.

Fox, William T. R. 1970. "The Causes of Peace and Conditions of War." *The Annals of the American Academy of Political and Social Science* 392(November):1–13.

Fudenberg, Drew & Jean Tirole. 1991. *Game Theory.* Cambridge: The M.I.T. Press.

Furasawa, Taiji & Quan Wen. 2001. "Unique Inefficient Perfect Equilibrium in a Stochastic Model of Bargaining with Complete Information." Department of Economics, Vanderbilt University, mimeo.

Goemans, Hein E. 2000. *War and Punishment: The Causes of War Termination and the First World War.* Princeton: Princeton University Press.

Goldberg, Samuel. 1986. *Introduction to Difference Equations.* New York: Dover. A reprint of the 1956 edition.

Grimmett, G. R. & D. R. Stirzaker. 1992. *Probability and Random Processes.* 2nd ed. Oxford: Oxford University Press.

Halperin, Morton H. 1970. "War Termination as a Problem in Civil-Military Relations." *The Annals of the American Academy of Political and Social Science* 392(November):86–95.

Handel, Michael I. 1978. "War Termination—A Critical Survey." Jerusalem Papers on Peace Problems, No. 24. The Leonard Davis Institute for International Relations, The Hebrew University of Jerusalem.

Kecskemeti, Paul. 1970. "Political Rationality in Ending War." *The Annals of the American Academy of Political and Social Science* 392(November):105–15.

Kim, Dong-Won. 2001. "Bargaining, Conflict, and War." Mimeo, Texas Tech University.

Manwaring, Max G. 1987. Limited War and Conflict Control. In *Conflict Termination and Military Strategy: Coercion, Persuasion, and War*, ed. Stephen J. Cimbala & Keith A. Dunn. Boulder: Westview Press pp. 59–76.

Maskin, Eric & Jean Tirole. 2001. "Markov Perfect Equilibrium I: Observable Actions." *Journal of Economic Theory* 100(October):191–219.

Merlo, Antonio & Charles Wilson. 1995. "A Stochastic Model of Sequential Bargaining with Complete Information." *Econometrica* 63(March):371–99.

Muthoo, Abhinay. 1999. *Bargaining Theory with Applications.* Cambridge: Cambridge University Press.

Norris, James R. 1997. *Markov Chains.* Cambridge: Cambridge University Press.

Oren, Nissan, ed. 1982. *Termination of Wars: Processes, Procedures, and Aftermaths.* Jerusalem: The Magnes Press.

Osborne, Martin J. & Ariel Rubinstein. 1990. *Bargaining and Markets.* San Diego: Academic Press.

PILLAR, PAUL R. 1983. *Negotiating Peace: War Termination as a Bargaining Process.* Princeton: Princeton University Press.

POWELL, ROBERT. 2001. "Bargaining While Fighting." Mimeo, University of California, Berkeley. October 7, 2001 version.

RUBINSTEIN, ARIEL. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50:97–110.

RUBINSTEIN, ARIEL. 1991. "Comments on the Interpretation of Game Theory." *Econometrica* 59(July):909–24.

SELTEN, REINHARD. 1975. "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." *International Journal of Game Theory* 4(1):25–55.

SIMON, CARL P. & LAWRENCE BLUME. 1994. *Mathematics for Economists.* New York: W. W. Norton & Company.

SMITH, ALASTAIR. 1998. "Fighting Battles, Winning Wars." *The Journal of Conflict Resolution* 42(June):301–20.

SMITH, ALASTAIR & ALAN C. STAM, III. 2001. "Bargaining Through Conflict." Paper presented at the Annual Meeting of the Public Choice Society, March 8-10, San Antonio, TX.

STÅHL, INGOLF. 1972. *Bargaining Theory.* Stockholm: Ekonomiska forskningsinstitutet vid Handelshögskolan i Stockholm.

TIROLE, JEAN. 1993. *The Theory of Industrial Organization.* Cambridge: The M.I.T. Press.

WAGNER, R. HARRISON. 2000. "Bargaining and War." *American Journal of Political Science* 44(July):469–84.

A The Closed Forms of W_k^i

In this section, I derive the closed form of the expected payoffs from fighting until the bitter end as specified in Definition 4.1. The functions W_k^i are second-order linear recurrence relations that we can solve by applying the usual techniques for solving difference equations. To simplify notation, let $a = \delta p$, $b = \delta(1-p)$, and divide through by a to obtain the standard form[10]

$$W_{k+2}^1 - \frac{1}{a}W_{k+1}^1 + \frac{b}{a}W_k^1 = \frac{-(1-\delta)b_1}{a} \tag{2}$$

The homogenous equation for (2) is

$$W_{k+2}^1 - \frac{1}{a}W_{k+1}^1 + \frac{b}{a}W_k^1 = 0 \tag{3}$$

We shall try a solution of the form $W_k^1 = m^k$. Substituting and dividing through by m^k yields the characteristic equation

$$m^2 - \frac{1}{a}m + \frac{b}{a} = 0$$

which has the solutions

$$m_1 = \frac{1+\sqrt{1-4ab}}{2a} \qquad m_2 = \frac{1-\sqrt{1-4ab}}{2a}$$

[10] I assume that both a and b are nonzero for otherwise the solution is trivial and uninteresting.

Note that $1 - 4ab = 1 - 4\delta^2 p(1-p) \geq 0$. The two roots are distinct as long as $1 - 4\delta^2 p(1-p) \neq 0$, or as long as, since $\delta \geq 0$,

$$\delta \neq \frac{1}{2}\sqrt{\frac{1}{p(1-p)}} \tag{4}$$

However since the function on the RHS attains a minimum of 1 at $p = \frac{1}{2}$, it follows that this condition is satisfied for all $\delta < 1$.[11]

The general solution to (3) has the form $W_k^1 = A_1 m_1^k + A_2 m_2^k$ where A_1 and A_2 are some arbitrary constants. We shall use the fact that the sum of the general solution to the homogenous equation and any particular solution to (2) yields the general solution we need (Goldberg 1986, Theorem 3.6). To find such a particular solution, we try a form $\overline{W}_k^1 = A_3$:

$$A_3 - \frac{1}{a}A_3 + \frac{b}{a}A_3 = A_3\left(\frac{a+b-1}{a}\right)$$

For this to be a solution, we must have

$$A_3\left(\frac{a+b-1}{a}\right) = \frac{-(1-\delta)b_1}{a}$$

which yields $A_3 = b_1$. Thus, the general solution to (2) has the form

$$W_k^1 = A_1 m_1^k + A_2 m_2^k + A_3 \tag{5}$$

Using the boundary conditions $W_0^1 = 0$ and $W_N^1 = \pi$, we can find the constant coefficients A_1 and A_2:

$$W_0^1 = A_1 m_1^0 + A_2 m_2^0 + b_1 = A_1 + A_2 + b_1 = 0$$
$$W_N^1 = A_1 m_1^N + A_2 m_2^N + b_1 = \pi$$

Thus, $-A_2 = A_1 + b_1$ and

$$A_1 = \frac{\pi - b_1\left(1 - m_2^N\right)}{m_1^N - m_2^N}$$

Substituting in (5) and rearranging terms yields the solution we want:

$$W_k^1 = \left[\pi - b_1\left(1 - m_2^N\right)\right]\left(\frac{m_1^k - m_2^k}{m_1^N - m_2^N}\right) + b_1\left(1 - m_2^k\right) \tag{6}$$

Hence, (6) is the complete solution to (2). W^1 is strictly increasing in k, p, b_1, and π.

To calculate the equivalent closed form for W_k^2, note that the standard form is very similar to (2):

$$W_{k+2}^2 - \frac{1}{a}W_{k+1}^2 + \frac{b}{a}W_k^2 = \frac{-(1-\delta)b_2}{a} \tag{7}$$

Since the homogenous equation is therefore the same, it follows that the roots of the characteristic equation are also the same. Hence, the general solution is

$$W_k^2 = B_1 m_1^k + B_2 m_2^k + B_3$$

[11] It is not difficult to verify that the derivative of the RHS is

$$\frac{2p-1}{4p^2(1-p)^2\sqrt{\frac{1}{p(1-p)}}}$$

and is zero at $\underline{p} = \frac{1}{2}$, negative for $p < \underline{p}$, and positive for $p > \underline{p}$. The second derivative, although harder to calculate, equals 4 at $\underline{p}$, and so the function attains a minimum there.

where B_1, B_2, and B_3 are some constants. For the particular solution to work, we must have

$$B_3\left(\frac{a+b-1}{a}\right)=\frac{-(1-\delta)b_2}{a}$$

which yields $B_3 = b_2$. Using the boundary conditions, we can find the other coefficients:

$$W_0^2 = B_1 m_1^0 + B_2 m_2^0 + b_2 = B_1 + B_2 + b_2 = \pi$$
$$W_N^2 = B_1 m_1^N + B_2 m_2^N + b_2 = 0$$

which yields $B_2 = \pi - b_2 - B_1$ and

$$B_1 = \frac{-\pi m_2^N - b_2(1-m_2^N)}{m_1^N - m_2^N}$$

Substituting and rearranging terms yields the solution we want:

$$W_k^2 = \left[-\pi m_2^N - b_2\left(1-m_2^N\right)\right]\left(\frac{m_1^k - m_2^k}{m_1^N - m_2^N}\right) + \pi m_2^k + b_2\left(1-m_2^k\right) \quad (8)$$

Hence, (8) is the complete solution to (7). W^2 is strictly decreasing in k, and in p, and strictly increasing in b_2, and in π.

B Calculating $R^i(x,t)$

In this section, I derive the expected payoffs from (x,t), an agreement on partition x reached in period t. The discussion below concerns player 1's payoffs, those for player 2 are computed in an analogous way.

B.1 Probability of Ending in the nth Period

We first need to know the probability of winning (losing) in exactly n steps when starting from k_0. Consider the probability that player 1 wins, that is, the probability of hitting N in the nth step when starting from k_0. For convenience, we shall call intermediate victories *right steps*, and intermediate defeats *left steps*, denoted by r and l respectively. Clearly, $n = r + l$. Also, since the total displacement to the right is $N - k_0$, it follows that $r - l = N - k_0$. Therefore, $r = \frac{1}{2}(n - k_0 + N)$. The probability associated with one path that visits N from k_0 in the nth step is $p^r(1-p)^l$, or

$$p^{\frac{1}{2}(n-k_0+N)}(1-p)^{\frac{1}{2}(n+k_0-N)} \quad (9)$$

Similarly, for paths that end in defeat for player 1 after n steps, the total displacement to the left is k_0, and so $l - r = k_0$. Then, $n = r + l \Rightarrow r = \frac{1}{2}(n-k_0)$, and $l = \frac{1}{2}(n+k_0)$, which yields the associated probability

$$p^{\frac{1}{2}(n-k_0)}(1-p)^{\frac{1}{2}(n+k_0)} \quad (10)$$

The probability of ending at N in n steps starting from k_0 is thus $p^r(1-p)^l$ times the number of paths that lead from k_0 to N. There are $\binom{n}{r}$ ways to choose paths with r right steps from n steps available. Unfortunately, some of these paths may be inadmissible because they may visit N earlier or they may visit 0. Since these are absorbing states, the game ends. We need to find a way to exclude these invalid paths and only count the remaining ones, which we shall call *valid*. To that end, we want to know how many paths of length n that start from k_0 visit N for the first time in their nth step without visiting 0 at any time. (We also want to know how many paths of length n visit 0 for the first time in their nth step without visiting N but this quantity can be calculated using the formula for the other paths.)

B.2 An Example of Counting Paths

Let $M_n(a,b) = \binom{n}{r}$, where $r = \frac{1}{2}(b - a + n)$, denote the number of paths which start from a, end in b, and have r rightward steps. Also, let $O_n(a,b) \le M_n(a,b)$ denote the number of valid paths.

To illustrate the logic that follows, we shall make use of an example with $N = 5$ and $k_0 = 2$. The shortest path is therefore $N - k_0 = 3$ right steps. Consider first $n = 3$. There is exactly $M_3(2,5) = \binom{3}{3} = 1$ path, and it is clearly valid. More generally, it is easy to see that the shortest path will always be valid. Thus, $O_3(2,5) = 1$.

Noting that all paths that end at N will have an odd number of steps in this case, consider next $n = 5$, which implies $r = 4$, $l = 1$. There are $M_5(k_0,N) = \binom{5}{4} = 5$ paths that start from k_0 and finish at N that are of length 5. We want to know how many of those are valid. Since there is no way to reach 0 with only one left step from k_0, we only need examine the paths that might reach N "too early." The only way this can happen is in paths that visit N in their third (right) step. From the preceding paragraph, there's exactly 1 such path. Conditional then on having reached N in three steps, how many ways are there that return to N in two more steps (one right and one left)? There are exactly $M_1(N,N) = \binom{2}{1} = 2$ such paths. Therefore, the number of invalid paths is $2 \times 1 = 2$, and so the total number of valid paths that start at k_0 and visit N in their 5th step for the first time, is $O_5(2,5) = 5 - 2 = 3$.

We confirm the result by inspection. The five possible paths are

$$(2,3,4,5,6,5) \quad (2,3,4,5,4,5) \quad (2,3,4,3,4,5)$$
$$(2,3,2,3,4,5) \quad (2,1,2,3,4,5)$$

Of these, the first two are clearly invalid because they both visit 5 for the first time in their third step instead of the fifth. Notice that once 5 is visited in the third step, there are only two ways to return to it, as obtained above.

Consider now $n = 7$. We have $r = \frac{1}{2}(5 - 2 + 7) = 5$, and so $M_7(2,5) = 21$. How many of these are invalid? There are two ways to reach N too early: in 3 steps and in 5 steps. We already know that there are $O_3(2,5) = 1$, and $O_5(2,5) = 3$ distinct valid such paths. Conditional on having reached N in 5 steps, there are $M_2(N,N) = 2$ ways to return to N in two more steps (here, we make use of the fact that reaching N in 5 steps requires 4 right and 1 left steps, and so we have $r-4 = 1$ right and $l-1 = 1$ left steps remaining). Thus, there are $O_5(2,5) \times M_2(N,N) = 3 \times 2 = 6$ invalid paths that visit N in their fifth step. Similarly, conditional on having visited N in the third step, there are $M_4(N,N) = 6$ ways to return to N in four more steps (here, we make use of the fact that reaching N in 3 steps requires 3 right and no left steps, and so we have $r-3 = 2$

right and $l - 0 = 2$ left steps remaining). Thus, there are $1 \times 6 = 6$ invalid paths that visit N in their third step.

This makes a total of 12 invalid paths that visit N too early. However, these do not include paths that might have also visited 0, and which would also be invalid. With two left steps, there is only one such path, which we must also exclude. This brings the total of invalid paths to 13. We therefore conclude that there are $O_7(2,5) = 21 - 13 = 8$ valid paths. The 21 possible paths are

$(2,3,4,5,6,7,6,5)$	$(2,1,2,3,4,5,4,5)$
$(2,3,4,5,6,5,6,5)$	$(2,3,4,3,4,5,6,5)$
$(2,3,4,5,4,5,6,5)$	$(2,3,2,3,4,5,6,5)$
$(2,3,4,5,6,5,4,5)$	$(2,1,2,3,4,5,6,5)$
$(2,3,4,5,4,5,4,5)$	$(2,3,4,3,4,5,4,5)$
$(2,3,4,5,4,3,4,5)$	$(2,3,2,3,4,5,4,5)$
$(2,1,0,1,2,3,4,5)$	$(2,1,2,3,4,3,4,5)$
	$(2,1,2,3,2,3,4,5)$
	$(2,1,2,1,2,3,4,5)$
	$(2,3,4,3,4,3,4,5)$
	$(2,3,4,3,2,3,4,5)$
	$(2,3,2,3,4,3,4,5)$
	$(2,3,2,3,2,3,4,5)$
	$(2,3,2,1,2,3,4,5)$

The realizations in the upper 6 entries first column are invalid because their first visit to N is in three steps and, as we calculated above, there are 6 such paths. The paths in the upper 6 entries in the second column are invalid because their first visit to N occurs in their fifth step. There are also 6 such paths. Finally, the path in the lower part of the first column is invalid because it visits 0. The eight remaining paths are the only valid ones that visit N for the first time in the seventh step without visiting 0.

Consider now $n = 9$, and so $r = 6$, $l = 3$, and $M_9(2,5) = 84$. How many of these are invalid? There are three ways to reach N too early: in 3, 5, or 7 steps, with $O_3(2,5) = 1$, $O_5(2,5) = 3$, and $O_7(2,5) = 8$ *distinct* valid paths respectively. Again, we condition on visiting N early to calculate the number of invalid paths. Conditional on having reached N in three steps, there are $M_6(N,N) = 20$ ways to return to N in six more steps (here we make use of the fact that reaching N for the first time in three steps requires three right and no left steps, and so we have $r - 3 = 3$ right and 3 left steps remaining). Thus, there are $1 \times 20 = 20$ invalid paths that visit N for the first time in their third step. Conditional on having reached N in five steps, there are $M_4(N,N) = 6$ ways to return to N in four more steps (here we make use of the fact that reaching N for the first time in five steps requires four right and one left step, and so we have $r - 4 = 2$ right and $l - 1 = 2$ left steps remaining). Thus, there are $3 \times 6 = 18$ invalid paths that visit N for the first time in their fifth step. Finally, conditional on having reached N in seven steps, there are $M_2(N,N) = 2$ ways to return to N in two more steps (again making use of the fact that reaching N for the first time in seven steps requires five right and two left steps, thus leaving $r - 5 = 1$ right and $l - 2 = 1$ left steps). Therefore, there are $8 \times 2 = 16$ invalid paths that visit N for the first time in their seventh step.

This makes a total of 54 paths that visit N too early using a valid path. How many of the remaining 30 paths visit 0? Note that all these paths visit N for the first time in their ninth step because the procedure above has already excluded paths that visit 0 and visit N too early. Therefore, we need to know how many paths that visit N for the first time in their ninth step also visit 0.

By the reflection principle, we can calculate the number of paths that reach N in their ninth step, and that visit 0 at some point. Since all such paths must visit 0 at some point for the first time, reflecting the segment prior to that visit about 0 yields a one-to-one correspondence, and so the number is the same as the number of paths that start from $-k_0$ and visit N in their nth step: $\binom{n}{\frac{1}{2}(n+N-(-k_0))}$. In our case, we have $\binom{9}{8} = 9$ such paths.

Thus, the number of *distinct valid* paths is $O_9(2,5) = 84 - 54 - 9 = 21$. It is important to note here that it is *not* the case that we are double counting paths. When we calculate the number of paths that visit 0, we are including paths that might also visit N before the 9th step. For example, one such path would be $\mathbf{s} = (2,1,0,1,2,3,4,5,4,5)$. So the question is: Since it is a path that visits N too early (in its seventh step), is it not excluded already by our procedure above? If this were the case, we would have a problem with subtracting the same path twice, once in the procedure that eliminates paths that visit N too early, and again in the procedure that eliminates paths that visit 0. Fortunately, this is not the case.

The reason the procedure that eliminates early hits did not remove the path $\mathbf{s}$ from the list is because the path that reaches N in seven steps is not valid, and so it is not included in the number $O_7(2,5)$. It is invalid because the relevant segment given by the sequence of steps $(2,1,0,1,2,3,4,5)$ visits 0, and so was eliminated during the previous step. The procedure that eliminates early hits for $n = 9$ *considers only valid paths* that reach N early. Thus, the recursive definition of the elimination procedure avoids double-counting paths. (Is this neat or what!)

There is only one problem left that we need to deal with. With a high n, there will be paths that are counted twice by the above exclusion algorithm: these are the paths that visit N early and then visit 0 before visiting N in the nth step. For example, for $n = 13$, the path $(2,3,4,5,4,3,2,1,0,1,2,3,4,5)$ does just that. Since the segment $(2,3,4,5)$ is a valid way to reach N in three steps, it will be included in the counting, which then multiplies that number by the number of ways to reach N from N in 10 more steps. That number will include the remainder of the path. However, since this includes 0, the path will also be counted by the procedure that excludes paths that visit 0, thus subtracting it twice.

The solution is to *exclude* all paths that visit 0 when counting the ways to revisit N in the remaining number of steps. This is straightforward: in the above example, the number of ways to revisit N in 10 steps is $\binom{10}{5} = 252$. By the reflection principle, the number of paths that visit 0 among these equals $\binom{10}{\frac{1}{2}(10+2N)} = 1$, and so there is precisely one path that is counted twice. The following section states the above algorithm generally and formally.

B.3 Counting Valid Paths: General Formula

As before, let $M_n(a,b) = \binom{n}{r}$ denote the number of paths of length n with $r = \frac{1}{2}(b - a + n)$ rightward steps such that the initial state is a, and the ending state is b. A path of length n which starts from k_0 and visits N for the first time in its nth step without visiting 0 is called *valid*. Let $O_n(k_0, N)$ denote the number of valid paths of length n, and define it recursively for all $n > 0$, where n even (odd) if $N - k_0$ is even (odd),

$$\begin{aligned} O_n(k_0,N) &= 0 \text{ if } n < N - k_0 \\ O_n(k_0,N) &= 1 \text{ if } n = N - k_0 \\ O_n(k_0,N) &= M_n(k_0,N) - M_n(-k_0,N) \\ &\quad - \sum_{i=1}^{r-1} O_{n-2i}(k_0,N)\Big[M_{2i}(N,N) - M_{2i}(-N,N)\Big] \end{aligned} \tag{11}$$

where $r = \frac{1}{2}(N - k_0 + n)$. $O_n(k_0, N)$ is nondecreasing in n whenever defined. That is, the number of valid paths from of length n cannot be smaller than the number of valid paths of shorter length.

To find the number of paths of length n that start from k_0 and visit 0 for the first time in their nth step without visiting N, note that by the reflection principle this is equivalent to finding the number of paths of length n that start from $N - k_0$ and visit N for the first time in their nth step without visiting 0 with the numbers of right and left steps reversed. To see this, consider the total displacement to the left, which is k_0, and which yields $r' = \frac{1}{2}(n - k_0)$ right steps, and $l' = \frac{1}{2}(n + k_0)$ left steps in these paths. In a path that starts from $N - k_0$ and reaches N in its nth step, the corresponding values are $r = \frac{1}{2}(n - (N - k_0) + N) = \frac{1}{2}(n + k_0)$, and $l = \frac{1}{2}(n - k_0)$, as we would intuitively expect.

Thus, let $O_n(k_0, 0)$ denote the number of valid paths (i.e. that are of length n, that visit 0 for the first time in their nth step, and that do not visit N) of length n, and define it as

$$O_n(k_0, 0) = O_n(N - k_0, N) \tag{12}$$

Note that if $N - k_0$ is even, then $O_n(k_0, N) = 0$ for all n odd, and if $N - k_0$ is odd, then $O_n(k_0, N) = 0$ for all n even. Similarly, if k_0 is even, then $O_n(k_0, 0) = 0$ for all n odd, and if k_0 is odd, then $O_n(k_0, 0) = 0$ for all n even. We can now turn to calculating the expected payoff from (x, t).

B.4 Calculating the Payoffs

Using (9) and (11), the probability of ending at N in period t, but not before, is

$$p_t(N) = O_t(k_0, N)p^{\frac{1}{2}(t-k_0+N)}(1 - p)^{\frac{1}{2}(t+k_0-N)}$$

Similarly, using (10) and (12), the probability of ending at 0 in period t, but not before, is

$$p_t(0) = O_t(N - k_0, N)p^{\frac{1}{2}(t-k_0)}(1 - p)^{\frac{1}{2}(t+k_0)}$$

The probability that the stochastic process ends exactly at some time t, but not before, is simply the probability that it visits either one of the absorbing states at that time:

$$P(t) = 1 - \big(1 - p_t(N)\big)\big(1 - p_t(0)\big) = p_t(N) + p_t(0)$$

where the second equality follows from the fact that reaching N and reaching 0 are mutually exclusive events. As we would expect, $P(t)$ is strictly decreasing for all $t \geq \min\{k_0, N - k_0\}$ and is 0 everywhere else. Finally, the probability that the stochastic process does not end until period t is the complement of the probability that it ends in some period prior to t. Since the events of ending in any period t but not before are mutually exclusive, we can simply sum over the their associated probabilities to find the probability that the process ends in some period prior to t. Thus, the probability of continuing until t is

$$1 - \sum_{\tau=0}^{t} P(\tau)$$

The summation ends with t, rather than $t - 1$ because if the game is not to end until period t starting from period 0, exactly t fights will occur. Since $\lim_{t\to\infty} \sum_{\tau=0}^{t} P(\tau) = 1$, the probability that the game will continue for a very large number of periods goes to zero.

Given that the process ends in period t, player 1's expected payoff is the lottery between π and 0, which is simply

$$\frac{\pi p_t(N)}{p_t(N) + p_t(0)} + \frac{(0)p_t(0)}{p_t(N) + p_t(0)} = \frac{\pi p_t(N)}{P(t)}$$

Player 1's expected payoff from (x, t) is the payoff from the distribution x agreed on at time t times the probability that the game lasts until t plus the expected payoff from victory and defeat in some period prior to t times the probability that this event occurs. Since $P(0) = 0$, we have

$$\begin{aligned} R^1(x,t) &= \left[1 - \sum_{\tau=0}^{t} P(\tau)\right]\Big[(1-\delta^t)b_1 + \delta^t x\Big] \\ &\quad + \sum_{\tau=0}^{t} P(\tau)\left[(1-\delta^\tau)b_1 + \delta^\tau \pi \frac{p_\tau(N)}{P(\tau)}\right] \end{aligned}$$

Note that $R^1(x, 0) = x$, and that for all $t < \min\{k_0, N - k_0\}$, $R^1(x,t) = (1-\delta^t)b_1 + \delta^t x$ because for such t, $P(t) = p_t(N) = p_t(0) = 0$. That is, the expected payoff in this case is simply the average of the per-period disagreement payoff and x, as in standard models with inside options. This is because the probability that the game reaches one of the absorbing states for such small t is zero.

Analogous calculations yield the expression for $R^2(x, t)$, player 2's payoff from partition x in period t.

$R^1(\cdot)$ is strictly increasing in x, and $R^2(\cdot)$ is strictly decreasing in x. Also, $\lim_{t\to\infty} R^i(\cdot, t) = W^i$. That is, in the limit, the expected payoff from (x, t) converges to the expected payoff from fighting to the end. To get some further insight into the behavior of the payoff functions, refer to Figure 3. Notice that the $R^i(x, t)$ is non-monotonic in t, where its behavior depends on the value of x, and which is not surprising. Consider player 1's expected payoff from (x, t) where x is a small partition (e.g. .1 in this example) and t is also small (e.g. 12). In this case, fighting to the end actually promises a better payoff because the probability of complete victory increases the expectation. On the other hand, when x is very large (e.g. .9), any delay is costly, and, for larger t, it is significantly so because the probability of defeat reduces the expected payoff.

C Proofs

Proof of Lemma 5.2. Since $\mathbf{w} \neq \mathbf{0}$, it is sufficient to establish that $\mathbf{A}^{-1}$ exists. $\mathbf{A}$ can be partitioned into four square submatrices:

$$\mathbf{A} = \begin{pmatrix} \mathbf{I} & \mathbf{M} \\ \mathbf{M} & \mathbf{I} \end{pmatrix}$$

where $\mathbf{I}$ is the identity matrix of size n and $\mathbf{M}$ is a $n \times n$ matrix with the following structure:

$$\begin{pmatrix} 0 & \delta p & 0 & \cdots\cdots\cdots\cdots\cdots & 0 \\ \delta(1-p) & 0 & \delta p & 0 & \cdots\cdots\cdots\cdots & 0 \\ 0 & \delta(1-p) & 0 & \delta p & 0 \quad \cdots\cdots\cdots & 0 \\ \vdots\ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots\cdots\cdots\cdots & 0 & \delta(1-p) & 0 & \delta p \\ 0 & \cdots\cdots\cdots\cdots\cdots & & 0 & \delta(1-p) & 0 \end{pmatrix}$$

In words, $\mathbf{M}$'s diagonal elements are 0, the lower off-diagonal elements are $\delta(1-p)$, the upper off-diagonal elements are δp, and the rest are all 0. Thus, each element of $\mathbf{M}$ is nonnegative and the sum of the entries in each column is less than 1.

Consider $\mathbf{M}^2$. The first and last diagonal elements are $\delta^2(1-p)p$ and the rest are $2\delta^2(1-p)p$. The lower off-diagonal elements, next to the 0s, are all $\delta^2(1-p)^2$, and the upper off-diagonal elements, again next to the 0s, are all $\delta^2 p^2$. Thus, each element of $\mathbf{M}^2$ is nonnegative. To see that the sum of entries in each column is less than 1, note that there can be at most five cases to consider: the first two columns, the last two columns, and the middle columns, which have the same sum. However, since the sum of elements in a middle column exceeds the sum in the other four cases because it always includes all corresponding terms plus an additional positive one, the required result follows from considering only this case, where the sum is given by

$$\delta^2 p^2 + 2\delta^2(1-p)p + \delta^2(1-p)^2 = \delta^2 < 1$$

By Theorem 8.13 in Simon & Blume (1994, p. 175), if $\mathbf{M}^2$ is a $n \times n$ matrix with the properties that $\mathbf{M}^2 \geq 0$ and the entries in each column sum to less than 1, then $(\mathbf{I} - \mathbf{M}^2)^{-1}$ exists and has only nonnegative entries. Thus, $\det(\mathbf{I} - \mathbf{M}^2) \neq 0$.

It can be shown that $\det \mathbf{A} = \det(\mathbf{I} - \mathbf{M}^2)$.[12] It then follows that $\det \mathbf{A} \neq 0$, and so the inverse exists. This establishes the result. Q.E.D.

Proof of Proposition 5.3. Consider player 1's proposal at some arbitrary time $2t$. Denote the current state by k. If player 1 follows the equilibrium strategy, its payoff is x_k^*. If it deviates and proposes some $x < x_k^*$, then player 2 accepts, which leaves player 1 worse off. Therefore, such deviation is not profitable. Suppose now player 1 offers some $x > x_k^*$, which player 2 always rejects. In that case, player 1's payoff is

$$(1-\delta)b_1 + \delta(\pi - p y_{k+1}^* - (1-p)y_{k-1}^*)$$

Suppose this deviation is profitable. From the definition of x_k^* from (1), this would then imply

$$\begin{aligned}(1-\delta)b_1 + \delta(\pi - p y_{k+1}^* - (1-p)y_{k-1}^*) \\ > \pi - (1-\delta)b_2 - \delta(p y_{k+1}^* + (1-p)y_{k-1}^*)\end{aligned}$$

or, after simplifying and rearranging terms,

$$\begin{aligned}(1-\delta)b_1 + \delta\pi > \pi - (1-\delta)b_2 \\ b_1 + b_2 > \pi\end{aligned}$$

which is a contradiction because $b_1 + b_2 \leq \pi$. This establishes that player 1 has no incentive to deviate by delaying agreement one period. Therefore, by the principle of optimality,[13] the proposal rule is optimal.

Consider now player 1's acceptance rule at some arbitrary time $2t+1$ and denote the current state by k. Suppose $y < y_k^*$, in which case player 1's payoff is $\pi - y$ if it accepts. If player 1 deviates and rejects, then its payoff is $(1-\delta)b_1 + \delta[p x_{k+1}^* + (1-p)x_{k-1}^*]$. Since $\pi - y > \pi - y_k^* = (1-\delta)b_1 + \delta[p x_{k+1}^* + (1-p)x_{k-1}^*]$, it follows that such deviation is not profitable. Suppose now that $y > y_k^*$, in which case player 1 should reject, getting a payoff of $(1-\delta)b_1 + \delta[p x_{k+1}^* + (1-p)x_{k-1}^*]$. Suppose player 1 deviates and accepts, in which case it gets $\pi - y < \pi - y_k^* = (1-\delta)b_1 + [p x_{k+1}^* + (1-p)x_{k-1}^*]$. Therefore, such deviation is not profitable. By the principle of optimality, the acceptance rule is optimal.

Thus, player 1's strategy is optimal in every possible subgame given that player 2 follows the strategy specified in the proposition. The proof for player 2 is equivalent, *mutatis mutandis*, and therefore the strategies constitute a stationary no-delay Markov perfect equilibrium.

By Lemma 5.2, the vector with proposals is unique, which implies that there exists at most one stationary no-delay MPE. Q.E.D.

[12] See Exercise 26.21 in Simon & Blume (1994, p. 735), which establishes the result for a partition into four arbitrary square matrices with the only requirement that the lower diagonal partition forms a nonsingular matrix. In our case, this partition is $\mathbf{I}$ with $\det \mathbf{I} = 1$.

[13] See Fudenberg & Tirole (1991, pp. 108-10) for a proof of this principle, which states that to verify whether a strategy profile in a stage-game is subgame perfect, it suffices to check whether players have an incentive to deviate from the strategy once and thereafter conform to its prescription.

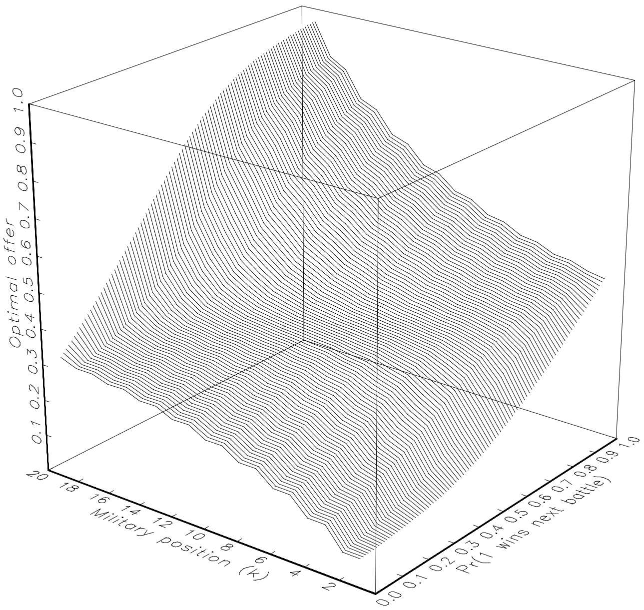

(a) $b_1 = 0; b_2 = .3; \delta = .9$

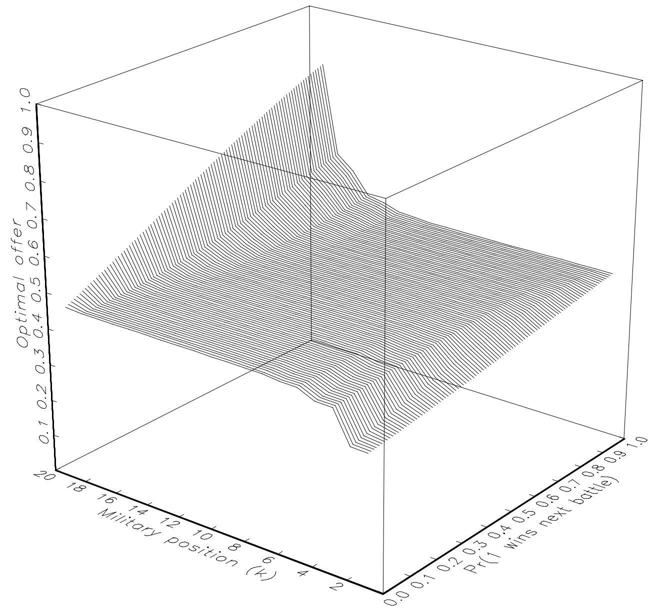

(b) $b_1 = 0; b_2 = .3; \delta = .5$

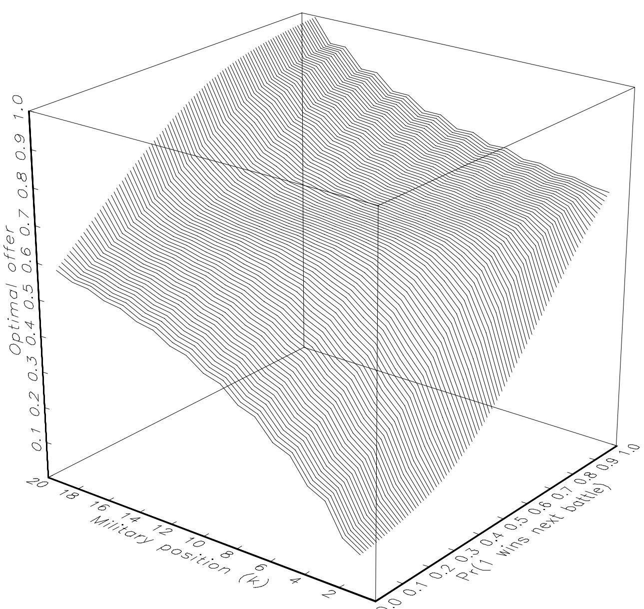

(c) $b_1 = .3; b_2 = 0; \delta = .9$

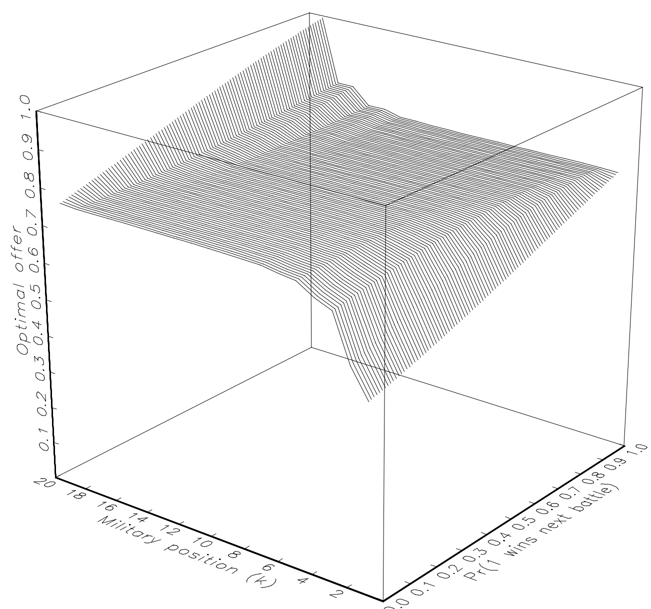

(d) $b_1 = .3; b_2 = 0; \delta = .5$

Figure 1: Player 1's equilibrium offers. $N = 20$, *other variables as shown.*

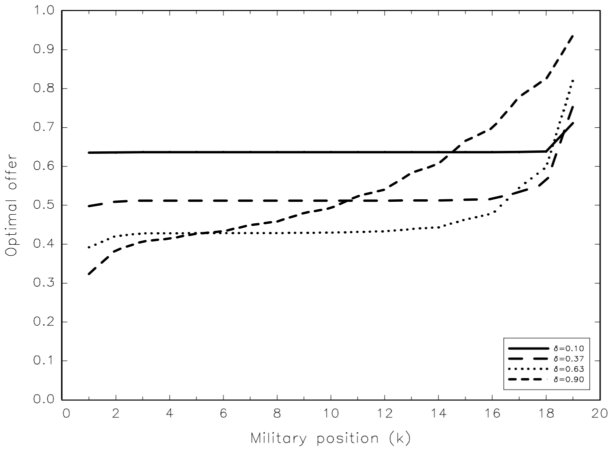

(a) $b_1 = 0; b_2 = .3$

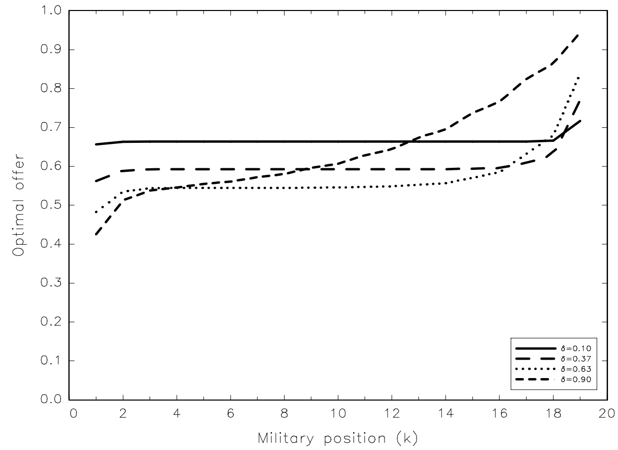

(b) $b_1 = .3; b_2 = .3$

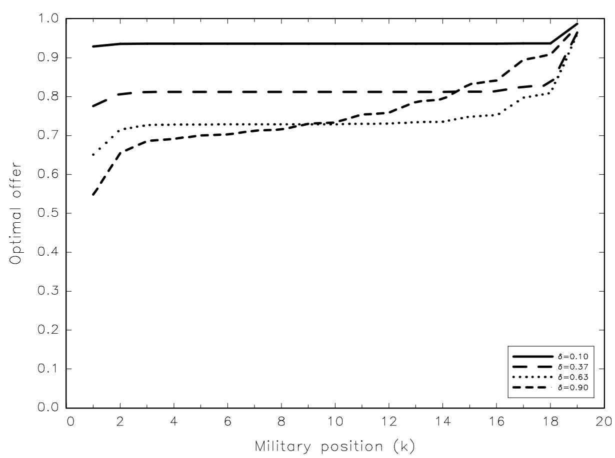

(c) $b_1 = 0; b_2 = .3$

Figure 2: Player 1's equilibrium offers for varying disagreement costs. *$N = 20$, $p = .8$, other variables as shown.*

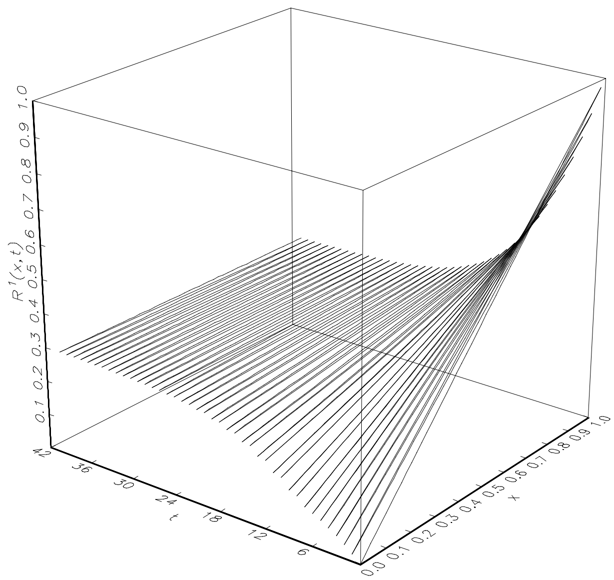

(a) $R^1(x,t)$

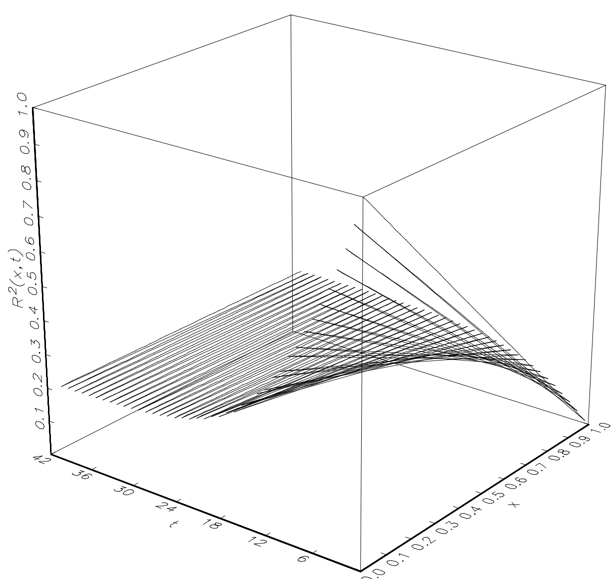

(b) $R^2(x,t)$

Figure 3: Expected payoffs from (x,t). The model parameters are $N = 20, k_0 = 12, p = .65, \delta = .9, \pi = 1, b_1 = .2, b_2 = .25$. Note that in this case, $W^1 \approx .3$ and $W^2 \approx .22$.