Bargaining and the Nature of War

ALASTAIR SMITH
Department of Politics
New York University
ALLAN C. STAM
Department of Government
Dartmouth College

A model of bargaining embedded within a random-walk model of warfare is developed. The conflict model contains aspects of both lottery-based and war-of-attrition models of conflict. Results show that future disputes are less likely to lead to armed conflict following long rather than short wars. Furthermore, should a subsequent dispute lead to armed conflict, the higher the cost and the longer the previous war, the shorter the conflict is likely to last.

Keywords: *warfare; bargaining; military; random-walk model; conflict*

Nations fight when they disagree about the nature of what a war between them would be like. In particular, when states disagree about who is likely to win or the costs of fighting, then they can fail to find a bargain that settles their differences (Blainey 1973). Nations therefore fight to resolve their differences of opinions. The act of waging war reveals information about the relative strengths of each side. As a war progresses, each side's beliefs about the likely outcome of continuing the war converge. Once the warring parties' beliefs have converged sufficiently, they can find a bargained solution to the conflict.

The contemporary modeling tradition in international relations has conventionally thought of war as the breakdown of bargaining. Early models portrayed war as a one-shot lottery. Empirically, however, wars vary vastly in their intensity and duration. Moreover, few wars are fought to complete disarmament of one side or the other. Instead, most wars end when both sides prefer a division of the stakes at issue to continued conflict. Rather than continuing to fight until one side eliminates the other, the two sides reach a bargain or negotiated settlement. Our model is one of a growing number of models that opens the black box of warfare and treats conflict as part of the bargaining process rather than the end of the bargaining process (Filson and Werner 2002, 2004; Morrow 1997; Powell 2004a, 2004b; Slantchev 2002, 2003, forthcoming;

AUTHORS' NOTE: The appendix with the proofs is available at http://www.yale.edu/unsy/jcr/jcrdata.htm.

JOURNAL OF CONFLICT RESOLUTION, Vol. 48 No. 6, December 2004 783-813
DOI: 10.1177/0022002704268026
© 2004 Sage Publications

Smith 1997, 1998; Smith and Stam 2003, forthcoming; Wagner 2000). Most of these models break down war from a single event to a series of interactions, or battles, during each of which nations choose actions such as the decision to continue prosecuting the war or whether to accept a bargained settlement to the conflict. Here we use Smith's (1998) random-walk model of conflict, in which nations A and B fight battles, during each of which either A captures a fort from B or B captures a fort from A. Here we embed a simple bargaining game within this model of war. The war continues until either one side exhausts its supply of forts, thereby being decisively defeated, or the nations agree to a bargain that ends hostilities. This framework operationalizes Wagner's (2000) concept of warfare and is also the basis of models by Filson and Werner (2002, 2004), Slantchev (2002, 2003, forthcoming), and Smith and Stam (2003, forthcoming).

Although our underlying model of war shares a similar structure to other models of war, we depart radically from the traditional approach in our treatment of informational beliefs. Although our informational structure resides within the standard Harsanyi (1967-1968, 1995) doctrine, we relax the conventional assumption of common priors. To our knowledge, within the political science literature, our model is the first game-theoretic model to dispense with the common-priors assumption. Because we believe that this is a significant technical innovation and a more accurate account of the empirical world, we use this introduction to contrast our approach, which we refer to as heterogeneous beliefs, to the standard private-information model of beliefs. Each approach has its limitations and pathologies for modeling the bargaining between potential belligerents (Morris 1995).

Harsanyi's (1967-1968, 1995) method of modeling beliefs is to partition the world into all its possible configurations, each of which is referred to as a state. In the context of war fighting, these states might refer to the probability, p, that nation A wins a battle. It might be the case that nation A is strong, and so p is high. Alternatively, it might be the case that nation B is the stronger, and so p is low. Both nations A and B have beliefs about the relative likelihood of these two states of the world. War occurs when nations disagree about what a war between them might be like—that is, which state of the world they are in.

If nations A and B were both certain that A was strong, they could divide the stakes accordingly. The nations could also agree if they both knew A was weak. Indeed, provided both nations agreed about the relative likelihood that nation A is strong, then agreement is straightforward (Fearon 1995). Unfortunately, reaching a peaceful negotiated settlement becomes more problematic if the nations have different assessments of the probability that A is strong. It is this disagreement that leads to war.

In casual language, the term *disagree* is readily understood. Yet within the Harsanyi (1967-1968, 1995) framework, disagreement can be conceptualized in a variety of ways, with the differences between them not trivial. This forces us to pause and ask the fundamental question: what does it mean to disagree? This is not simply a question of semantics. The answer alters what aspects of strategic behavior lead to war. Within the private-information conception of disagreement, credibility and the incentives to bluff dominate our understanding of strategic incentives and behavior. In contrast, the heterogeneous-beliefs approach emphasizes learning and each actor's underlying conceptual beliefs about the nature of the world. For instance, in political science, we might expect academics with Marxist, realist, and liberal perspectives to interpret the same political events very differently, even when given the same information and awareness of other alternative perspectives.

Most game-theoretic applications, by convention, assume that both nations initially have identical prior beliefs. Disagreement between the nations occurs when one nation receives information that the other does not. This information is typically called *private information* because it is not public or common knowledge to both nations. A simple example might be that nation A learns its true strength, and hence the state of world, whereas nation B does not receive this information. The information A learns is typically referred to as A's "type." The common prior beliefs assumption means that prior to nation A learning whether it is strong, both nations A and B placed the same probability on nation A being strong. Although A's information changes A's belief about its strength, B has not received any additional information beyond its initial, or prior, belief. Because A and B shared this prior belief, A knows what B believes, and B knows that A knows what B believes. Disagreement is created by the differential information each nation receives.

The private-information setup provides a very flexible modeling platform. Because nation A conditions its actions on the information it has received (its type), its actions may be potentially informative for nation B, in that they indicate the information observed by A but not by B. For instance, if A saw information that it was strong, then it might behave in a more belligerent manner than if it learned that it was weak. In this setting, nation B learns indirectly about the information A saw by observing the changes in A's level of belligerence. In this way, B updates its beliefs about the state of the world. In many cases, of course, B cannot completely discern A's type because A has incentives to bluff. That is, because weak types of nation A know that they will be exploited by nation B, they often try to mimic strong types of nation A to avoid being identified as a weak type. Recent efforts to understand how nations can signal their type has led to domestic political audience cost arguments, in which national leaders pay political costs for escalating crises and then backing down (Schultz 2001).

Although a few general results have been established for models of crisis bargaining and conflict, predictions are often extremely sensitive to precise details of the model. For example, Powell (2004a) and Lewis and Schultz (2003) show that subtle changes in the bargaining protocol or equilibrium-solution concept drastically alter the equilibrium behavior we expect to observe. In these models, although the act of fighting battles reveals information about the underlying balance of power between nations A and B (i.e., the state of the world), the informational content of such objective outcomes or signals is often dwarfed by the information revealed through the nations' bargaining strategies. Unfortunately, the revelation of this later information is quite sensitive to the precise modeling of the bargaining protocol (Powell 2004a). Because the inherent anarchy in the international system prevents commitment to a specific protocol, the private-information approach often fails to give us a clear general picture of how fighting resolves disagreements between nations (Filson and Werner 2004; Schultz 2001).

The convention in the private-information approach is to assume common prior beliefs; that is, before receiving private information, all actors share the same identical beliefs. The heterogeneous-beliefs approach that we employ instead focuses on underlying differences in how actors think the world works. In terms of modeling, this means we assume that the actors possess noncommon priors. Although the meaning of the term *common prior* is readily apparent to modelers, to those unfamiliar with probability models, it is a notion that deserves explanation. The common prior assumes that, initially, each player agrees with the other about the relative likelihood of each possible state in the distribution of states of the world. In the context of war, this means that both A and B think the probability that nation A is strong is the same.

The private-information approach creates differences in beliefs between the actors by giving them differential access to information. For example, only nation A might know the cohesion of units within its army. The common-priors assumption means that if A and B both see the same information, then beliefs about the state of the world converge (Aumann 1976). That is, if both nations A and B thought that cohesion within army units was important in determining military strength, then if both A and B observe the same information about A's military cohesion, then A and B would share the beliefs about the probability of A winning.

Unfortunately, common priors is a poor representation of the real world. Consider, for example, jury trials. Because all jurors receive the same information at the trial, if we assume that all jurors are truth seeking because they share the same common prior—that the defendant is innocent until proven guilty—and all the jurors see the same information during the trial, then every jury decision should be unanimous. We know, however, that they are not.

We argue that, even absent private information, actors can have different, or heterogeneous, beliefs that lead the actors into disagreement. One reason that beliefs fail to converge in the presence of the same information is that actors have different underlying theoretical perspectives. Although nations A and B might both agree about the cohesion of A's army, they might disagree about how important cohesion is in terms of measuring strength. As we discuss later, prior to the Seven Weeks War, Austria regarded unit cohesion as paramount. In contrast, Prussia thought it of much less importance relative to firepower. Although both nations observed the American Civil War, they drew very different conclusions from it with regard to their relative strengths.

To model disagreement between nations in terms of heterogeneous beliefs, we use the Harsanyi (1967-1968, 1995) setup but relax the assumption of common prior beliefs and allow actors to begin the game with different views of the world, or prior beliefs. Although to many game theorists, this is a controversial modeling approach, scholars such as Boge and Eisele (1979), Mertens and Zamir (1985), Brandenburger and Dekel (1993), Nyarko (1990, 1997), Morris (1995), and Gul (1998) construct logically consistent information models that allow for both common and noncommon prior beliefs. Within this general setting, the case of common prior knowledge is a special case—a nongeneric one, according to Nyarko (1991). Thus, although the common prior assumption fits within Harsanyi's more general framework, there is no logical basis for insisting on it. Morris (1995) provides a comprehensive summary of the arguments for and against the common knowledge assumption.[1] Perhaps the most persuasive argument Morris gives in favor of the common prior assumption over a heterogeneous-beliefs approach is the practical utility of the common prior assumption. Although infinite hierarchies of beliefs, of the form of "you know that I know that you know. . . ," can be consistently constructed with a variety of initial beliefs, without some kind of restriction on the actors' priors, they are often too unwieldy or admit all forms of behavior as equilibrium behavior.

We assume that although actors have different beliefs, what these beliefs are is common knowledge to all actors. That is, although nation A has different beliefs than nation B, nation A knows what nation B believes and vice versa. As was the case in 1866, Austria thought it was stronger than Prussia, and Prussia thought it was stronger than Austria. Each was aware of the other's view; they just thought each other was wrong. In essence, nations "agree to disagree." If the common-priors assumption is maintained, then "agreeing to disagree" is impossible. If Austria knew that Prussia thought itself stronger than Austria, then the Austrians would conclude that Prussia had seen information indicating the superior strength of the Prussian army. Knowing Prussia had this information, Austria would downgrade its assessment of its relative strength. Similarly, Prussia would downgrade its assessment of relative strength in light of Austria's beliefs. Unless the nations fundamentally disagree about what constitutes military strength, then beliefs converge: nations cannot agree to disagree if they have common priors.

The restriction that each side knows the beliefs of the other creates a readily manageable framework within which we explore how war reveals information to resolve the disagreements between A and B. Although we do not, in general, advocate replacing the traditional Harsanyi (1967-1968, 1995) approach, in this particular context, the heterogeneous-beliefs approach provides a more tractable framework and a more empirically accurate representation of the process of bargaining and war.

We proceed as follows: we introduce our underlying model of conflict, the random walk. We explain the specific informational structure of our heterogeneous-beliefs approach within the context of this model. These beliefs shape nations' expectations of the duration of conflict and which nation is likely to be the eventual winner if the war is fought to a decisive conclusion. Fighting until a decisive victory might be thought of as von Clausewitz's (1976) concept of absolute war. The wars with bargained settlements are more akin to von Clausewitz's limited wars, or wars in reality. A key feature of our model is that as nations fight battles and capture forts from each other, both nations learn common information about the nature of warfare between them. As more and more information is revealed, the nations' beliefs converge. Wars are fought until either one side decisively defeats the other or until each side's beliefs converge sufficiently that they can agree to a settlement. By characterizing the subgame-perfect equilibrium of the bargaining and conflict game, we examine how beliefs change during the course of the game. The model shows how war resolves differences between states' beliefs. Yet, such a process does not explain how beliefs diverge in the first place. We examine the consequences of assuming that nations have different underly-

1. See also Bernheim (1986) and Varian (1989).

ing theoretical conceptions about what factors determine the probability of victory in battle, p. In this setting, although battles between A and B serve to resolve the disagreement over p (the probability of victory in battle), observations of conflict between other pairs of nations can lead to a divergence of beliefs for nations A and B. We illustrate this differential learning process by examining how Prussia and Austria both learned different messages from the American Civil War. As a result, both Germanic powers thought themselves stronger than their opponent. These beliefs persisted despite each knowing of the other's differing beliefs. The Seven Weeks War revealed the relative veracity of each side's beliefs. Our results suggest relationships between the course of one war and the prospects of future conflicts, as well as the likely duration of any such future conflicts. We conclude by exploring these relations.

A RANDOM-WALK MODEL OF WAR WITH BARGAINING: THE MODEL

We begin by assuming the existence of two nations that are in dispute over the division of some prize, worth V to each player. We further assume that disputes arise because nations disagree about p, their likelihood of winning any individual battle. We defer for the moment the discussion of how nations' beliefs about p differ and, for now, simply assert that such differences exist. Differences in beliefs about p, the probability of winning an individual battle, mean nations may disagree about how they should split the prize. If nations fail to reach agreement, then they fight. As an alternative modeling strategy, one might model disagreement over the costliness of fighting (Smith and Stam, forthcoming).

In our setup, we treat war as a stochastic process akin to the gambler's ruin. The gambler's ruin is a well-known stochastic model in which two players gamble a fixed-size wager in each period. For the sake of illustration, suppose there are N forts and that, initially (at time $t = 0$), nation A possesses X_0 of them. Each time the nations fight, either A captures a fort from B, or B captures a fort from A. These events occur with probability p and $q = 1 - p$, respectively. If A wins the first battle, then the distribution of forts becomes $X_1 = X_0 + 1$; similarly, if A loses, then $X_1 = X_0 - 1$. We refer to the distribution of forts as the "state" variable. The war continues either until the nations agree to a division of the prize or until one nation decisively defeats the other. Decisive victory for nation A, $X_t = N$, means A has captured all the forts and B can no longer resist. At this point, nation A can impose any settlement on B. Similarly, the state $X_t = 0$ represents A being unable to prosecute the war. The $(N + 1) \times (N + 1)$ transition matrix P describes movements between states. For example, the element ij is the probability of moving from state $i - 1$ to state $j - 1$.

For convenience, we use the heuristic device of forts to describe the state variable. Modern battles typically do not revolve around forts or castles. Forts in a modern war of coercion might better represent units of resolve or strategically important pieces of land, for instance.[2] We stick to the term *forts* because its meaning is intuitive for many,

2. See Filson and Werner (2002) for a discussion of resolve as a unit in attrition warfare.

but in reality, we are thinking of any particular form of military resource that is necessary for the two nations to continue to fight and that can trade hands.[3]

Within the context of this stochastic model of war, we embed a simple bargaining game. In each round of the game, nations choose to either fight another battle or agree about the division of the prize. Although the characteristics of the bargaining protocol that we employ affect the precise deals we would expect the nations to reach, our general results are not dependent on the choice of any particular bargaining setup. Therefore, in the interests of space and transparency, we use a very simple take-it-or-leave-it bargaining setup. In the appendix (on the Web site), we analyze a more complex alternating-offers bargaining framework and show that it produces similar results to the one-sided bargaining used here.

If nations fail to reach a bargaining agreement over the division of the prize, then they fight a battle in which A captures a fort from B, or B captures a fort from A. Following a battle, the nations bargain again to try to reach an agreeable division of the prize and hence a settlement to the war. Again, if they fail to agree, they fight another battle. The game continues until they reach a negotiated settlement or one side decisively defeats the other.

The value of winning the prize is V, and nations pay a per period cost of K for each period of fighting. If, in period t, nation A proposes an agreement $b_t \in [0, V]$, which is accepted by B, then A's payoff is $V - b_t - Kt$, where $V - b_t$ represents A's share of the prize, and Kt is the cost of fighting for the t periods before reaching an agreement. B's corresponding payoff is $b_t - Kt$.[4] By a negotiated settlement, we mean the nations agree to stop fighting before reaching absorption (state 0 or N), where one side or the other captures all the forts. In this context, surrendering and completely relinquishing the prize to the other side is a negotiated settlement with divisions V and 0.

STAGE GAME

Consider the subgame starting in period t in state X_t and the simple take-it-or-leave-it, one-sided bargaining protocol.

(1) A proposes a division of the prize, $b_t \in [0, V]$, where b_t represents the amount of the prize A is willing to give B, with A retaining the remaining portion of the prize.

(2a) If B agrees to the proposed division, then the two sides divide the prize according to the proposed division, and the game ends. A's payoff is $V - b_t - Kt$, which represents A receiving its proportion of the prize $(V - b_t)$ less the cost of the conflict accumulated over the previous t periods. B's payoff is $b_t - tK$, which represents its share of the prize (b_t) less the cost of t periods of fighting (tK).[5]

(2b) If B rejects A's offer, then the two sides fight a battle. A wins the battle with probability p and loses with probability $q = 1 - p$. If A wins, then the state variable moves to $X_{t+1} = X_t + 1$, and if A loses, then the state moves to $X_{t+1} = X_t - 1$.

3. Elsewhere (Smith and Stam 2003; Smith 1997), we model p as a function of X.

4. Payoffs are not discounted. This presents a potential problem in that payoffs are unbounded. Fortunately, as we show next, from any given point in the game, the game ends in a finite expected time for any strategy by the players.

5. The addition of risk propensity is straightforward but ignored here.

If $X_{t+1} = N$, then A objectively wins the war, decisively defeating B. At this point, the game ends, with A receiving a payoff of $V - (t + 1)K$. This payoff reflects winning the prize (worth V) but paying the cost of $t + 1$ battles (periods of conflict). B's payoff is $0 - (t + 1)K$. If $X_{t+1} = 0$, then B decisively wins the war, defeating A. This ends the game, with associated payoffs for A and B of $0 - (t + 1)K$ and $V - (t + 1)K$. All other war outcomes are subjective, with each side's perception of victory or defeat conditioned on its ex ante expectations of how it ought to do in the war.

BELIEFS AND INFORMATIONAL STRUCTURE

We assume the structure and payoffs of the game to be common knowledge. The game begins with the assumption that nations A and B have different prior beliefs about the distribution of p, the probability that A wins the next battle. However, as discussed above, in contrast to the conventional signaling game setup, we assume that both A's and B's prior beliefs about the distribution of p are common knowledge.[6] In other words, A knows what B believes about the distribution of p yet sees the same information; A comes to a different conclusion than B about the value of p, the probability that A will win the next battle if they fight. Similarly, B knows that A's beliefs are different from its own. In effect, holding heterogeneous beliefs, nations A and B agree to disagree.

After exploring our model of bargaining and conflict, we will return to the question of how a divergence of the actors' prior beliefs can arise in the first place and the implications that flow from this modeling choice compared to the more standard setup driven by private information. Next, we describe the distributions of the actors' beliefs and the learning process that takes place during a war.

DISTRIBUTION OF BELIEFS

We let $F(p)$ represent the distribution (with associated density $f(p)$) of a nation's beliefs over p. Although not critical to our argument, for the purposes of constructing examples, we shall assume that nations' beliefs about the value of p follow the beta distribution. The beta distribution has two parameters, α and β, and its density function over the interval [0, 1] is

$$f(p;\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{(\alpha-1)} q^{(\beta-1)},$$

where $q = 1 - p$ and $\Gamma(\alpha) = \int_0^\infty e^{-x} x^{\alpha-1}\,dx$. The beta distribution has mean

$$\frac{\alpha}{\alpha+\beta}, \text{ variance } \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}, \text{ and mode } \frac{\alpha-1}{\alpha+\beta-2} \text{ (if } \alpha, \beta > 1).$$

6. See Slantchev (2003) for a similar model that focuses on private information, as well as Lewis and Schultz (2003) for one that relates the sensitivity of those models to assumptions about equilibrium solution concepts.

Nations A and B have different prior beliefs about p, represented in the case of the beta distribution through different α and β parameters. In particular, A's prior beliefs about the distribution of p are represented by the parameters α_{A0} and β_{A0}. B's prior beliefs are α_{B0} and β_{B0}.

As nations observe the outcomes of the war's battles, they update their beliefs about the distribution of p via the Bayes rule. If at time t, A has experienced w_t victories and l_t defeats, then A's posterior beliefs are that p is beta distributed with parameters $\alpha_{At} = \alpha_{A0} + w_t$ and $\beta_{At} = \beta_{A0} + l_t$. Because

$$w_t = \frac{t + X_t - X_0}{2} \text{ and } l_t = \frac{t - X_t + X_0}{2},$$

A and B's beliefs can be expressed in terms of the state variable and time period.

To illustrate our results, we use the running numerical example of $N = 10, X_0 = 5, V = 30, K = 1, \alpha_{A0} = 4, \beta_{A0} = 2, \alpha_{B0} = 2$, and $\beta_{B0} = 4$. For our example, A initially believes the expected value of p is 2/3, whereas B's expectations are that $p = 1/3$.

EXPECTATIONS OF WAR OUTCOME AND DURATION

Starting from state X_τ in time period $\tau \geq 0$, we define the resolution time, $R(X_\tau) = \min\{(t - \tau) \geq 0 | X_t = 0 \text{ or } X_t = N\}$, as the time taken until either nation A or nation B attains complete victory (states $X_t = N$, 0, respectively). We define the probability that nation A will eventually lose the war if it is fought to the end while starting in state y, given beliefs α, β, as $L(\alpha, \beta, y) = \Pr\{X_R = 0 | X_\tau = y\}$. With complementary probability, A will emerge the eventual winner: $W(\alpha, \beta, y) = \Pr\{X_R = N | X_\tau = y\} = 1 - L(\alpha, \beta, y)$.

Finally, we define $D(\alpha, \beta, y) = E[R(y)]$ as the expected number of battles before either A or B will win the war, given that we start in state y with beliefs α and β.

To calculate $L(\alpha, \beta, y)$, $W(\alpha, \beta, y)$, and $D(\alpha, \beta, y)$, we use some results from the stochastic processes literature. In particular, if the transition probabilities are p and $q = 1 - p$ and the current state is y, then the probability of reaching the absorbing state 0 before reaching the absorbing state N is

$$\lambda(p, y) = \begin{cases} \dfrac{(q/p)^y - (q/p)^N}{1 - (q/p)^N} & \text{if } p \neq 1/2 \\ \dfrac{(N - y)}{N} & \text{if } p = 1/2 \end{cases}.$$

The probability of reaching state N first is simply $1 - \lambda(p, y)$. The expected number of periods until reaching either state 0 or N is

$$E[R | p, y] = \begin{cases} \dfrac{1}{q - p}\left[y - N\left(\dfrac{1 - (\frac{q}{p})^y}{1 - (\frac{q}{p})^N} \right) \right] & \text{if } p \neq 1/2 \\ y(N - y) & \text{if } p = 1/2 \end{cases}.$$

B's belief that it will be the eventual winner should the war be fought out until a decisive winner emerges.

		t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7	t = 8	t = 9
	0						1		1		1
	1					1		0.99		0.99	
	2				0.98		0.97		0.96		0.95
	3			0.96		0.94		0.92		0.9	
	4		0.91		0.87		0.84		0.82		0.8
State, X_t	5	0.82		0.77		0.73		0.71		0.69	
	6		0.62		0.59		0.56		0.55		0.53
	7			0.4		0.39		0.39		0.38	
	8				0.22		0.22		0.23		0.23
	9					0.09		0.09		0.1	
	10						0		0		0

B's belief about the expected time until the war ends should the war be fought out until a decisive winner emerges.

		t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7	t = 8	t = 9
	0						0		0		0
	1					8.53		10.42		11.97	
	2				9.24		11.29		12.9		14.18
	3			10.15		12.32		13.92		15.16	
	4		11.26		13.46		14.99		16.12		17
State, X_t	5	12.52		14.6		15.98		16.97		17.74	
	6		15.46		16.65		17.53		18.2		18.74
	7			16.65		17.53		18.2		18.74	
	8				16.65		17.53		18.2		18.74
	9					15.98		16.97		17.74	
	10						0		0		0

Figure 1: B's Beliefs That It Will Be the Eventual Winner ($L(X_t, t)$) and the Expected Length of Conflict ($D(X_t, t)$)

Nations A and B are uncertain as to the true value of p. However, given beliefs α and β, the following integrals characterize $L(\alpha, \beta, y)$, $W(\alpha, \beta, y)$, and $D(\alpha, \beta, y)$:

$$L(\alpha, \beta, y) = \int_0^1 \lambda(p, y) f(p; \alpha, \beta) dp,$$

$$W(\alpha, \beta, y) = \int_0^1 (1 - \lambda(p, y)) f(p; \alpha, \beta) dp,$$

$$D(\alpha, \beta, y) = \int_0^1 E[R | p, y] f(p; \alpha, \beta) dp.$$

Figure 1 shows how

$$L\left(\alpha_{B0} + \frac{t + X_t - X_0}{2}, \beta_{B0} + \frac{t - X_t + X_0}{2}, X_t\right) \text{ and}$$
$$D\left(\alpha_{B0} + \frac{t + X_t - X_0}{2}, \beta_{B0} + \frac{t - X_t + X_0}{2}, X_t\right)$$

change as the game proceeds using our numerical example. The corresponding numbers for nation A can be calculated by inverting the table because the chosen parameters are symmetric.

In the initial state $X_0 = 5$, nation B's leaders believe they have a 0.82 chance of emerging as the decisive winner if the war is fought until one side eliminates the other.

The lower half of Figure 1 shows that B's leaders believe that such a war would take, on average, a little more than 12 periods to fight to the end. If B wins the first battle ($X_1 = 4$), then B's leaders subsequently update their beliefs and then believe their chance of eventual victory is 0.91, with the war expected to last an additional 11 periods. Winning the first battle improves B's prospects of victory through two mechanisms. First, B now only needs to win a net of four battles before A secures six net wins. Second, having won one battle, B's beliefs about p shift such that, in expectation, B's leaders believe that A has only a 2/7 chance of winning the next battle. If B loses the first battle, then its leaders believe that the expected value of winning the next battle is 3/7. The resultant shift in beliefs, as well as the fact that B must now win six net battles before A wins four net battles, means that B's overall probability of winning the war drops to 0.62, with the war expected to last 15 more periods. Figure 1 shows how the probability of eventual victory in the war and the expected additional length of a war change as the war is fought, where $L(\alpha, \beta, y)$ represents B's beliefs about its eventual probability of victory, and $D(\alpha, \beta, y)$ represents B's beliefs about the expected duration of the conflict.

THE VALUE OF FIGHTING

To estimate whether holding out for total victory is worth the extra cost of fighting the additional battles or whether the nations will prefer to settle for a bargain short of total victory, we next calculate the additional expected value of fighting the war until an eventual winner emerges. Suppose at time t and state X_t, nations A and B calculate the expected value of fighting the war to an absorbing state, less any cost from fighting already accumulated. Given X_t and t, A's and B's beliefs are

$$\left(\alpha_{A0} + \frac{t + X_t - X_0}{2} \text{ and } \beta_{A0} + \frac{t - X_t + X_0}{2}\right) \text{and} \left(\alpha_{B0} + \frac{t + X_t - X_0}{2} \text{ and } \beta_{B0} + \frac{t - X_t + X_0}{2}\right),$$

respectively.

Hence,

$$\begin{aligned} E[U_a(\mathit{fight-indefinitely} | X_t, t)] &= C_A(X_t, t) \\ &= VW\left(\alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t + X_t + X_0}{2}, X_t\right) \\ &\quad - KD\left(\alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t - X_t + X_0}{2}, X_t\right). \end{aligned}$$

$$\begin{aligned} E[U_b(\mathit{fight-indefinitely} | X_t, t)] &= C_B(X_t, t) \\ &= VL\left(\alpha_{B0} + \frac{t + X_t - X_0}{2}, \beta_{B0} + \frac{t + X_t + X_0}{2}, X_t\right) \\ &\quad - KD\left(\alpha_{B0} + \frac{t + X_t - X_0}{2}, \beta_{B0} + \frac{t - X_t + X_0}{2}, X_t\right). \end{aligned}$$

These expressions represent the overall probability of victory multiplied by the value of victory (V), less the expected duration of the conflict multiplied by the per period cost (K). These calculations are made from the perspective of each individual nation for the subgame starting at X_t. $C_B(X_t, t)$ is B's estimate of what B can expect to receive by fighting the war to a conclusion. However, we can also calculate B's expected value of fighting the war from A's perspective:

$$\widehat{C_B^A(X_t,t)} = VL\left(\alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t - X_t + X_0}{2}, X_t\right)$$
$$- KD\left(\alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t - X_t + X_0}{2}, X_t\right).$$

By dominance, B never accepts any agreement $b_t < C_B(X_t, t)$ because B always expects to do better by continuing to fight. We define $M(X_t, t) = (V - \max\{0, C_B(X_t, t)\})$. The max is included to reflect that all deals must be on the interval $[0, V]$. Hence, even if $C_B(X_t, t)$ is negative, A cannot offer B a deal worth less than 0. Also by dominance, A never proposes an agreement, b_t (which B would accept), that is worth less than continued fighting: $V - b_t > C_A(X_t, t)$. From this, our first result follows directly:

Proposition 1: A necessary condition for a negotiated settlement is $C_A(X_t, t) \leq M(X_t, t)$. This implies that fighting always continues if $C_B(X_t, t) + C_A(X_t, t) > V$.

Proposition 1 states formally the obvious proposition that if there is no division of the prize that both sides prefer to continued fighting, then a negotiated settlement is impossible. This result holds for any bargaining arrangement in which settlements require the consent of each party and not simply the one-sided, take-it-or-leave-it offers we examine here. However, there is an aspect to this proposition that is not obvious. There are divisions of the prize that both sides prefer to indefinite fighting, but this may be insufficient to ensure a negotiated settlement, as we shall see in our numerical example.

AN INFORMAL DESCRIPTION OF THE EQUILIBRIUM

Before proceeding to the remaining formal propositions, we pause to describe more nuanced play of the game. In equilibrium, A never offers B any deal larger than $\max\{0, C_B(X_t, t)\}$. Because B is never offered any larger deals, B can never expect to receive any payoff greater than $\max\{0, C_B(X_t, t)\}$ in any future period. Hence, B accepts any offer $b_t \geq C_B(X_t, t)$ but rejects offers $b_t < C_B(X_t, t)$. Given B's acceptance strategy, nation A can end the war with a negotiated settlement by offering nation B the deal $b_t = C_B(X_t, t)$. However, nation A's leaders might prefer to continue the war because they believe that doing so enhances their bargaining position in future periods. $C_B(X_t, t)$ and $\widehat{C_B^A(X_t,t)}$ provide a useful heuristic to explain the logic for postponing an agreement today for a more costly but, it is hoped, better agreement tomorrow.

Given B's beliefs about the world, $C_B(X_t, t)$ represents the minimum deal B would accept rather than fight the war to its conclusion. In contrast, given A's beliefs about the world, A believes the minimum deal B should accept is $\widehat{C_B^A(X_t,t)}$. If $\widehat{C_B^A(X_t,t)} > C_B(X_t,t)$, then A's leaders will happily offer B a smaller deal than they believe B should be prepared to settle for. Yet if $\widehat{C_B^A(X_t,t)} < C_B(X_t,t)$, then, from A's perspective, B is overly optimistic about its chances and, as a result, is demanding more than it really has a right to, at least given A's view of the world.

In this latter case, A must trade off two alternatives. First, A could offer B a greater share of the prize than A believes B deserves and conclude an agreement today; that is, $b_t = C_B(X_t,t) > \widehat{C_B^A(X_t,t)}$. Second, A could continue the war. By doing so, A attempts to convince B that A's view of the world is more accurate than B's. In this type of situation, given nation A's beliefs, its leaders believe that, on average, they will get a better deal tomorrow than they could obtain today because they believe the outcome of future battles will both reinforce their beliefs and also bring B around to A's way of thinking. Nevertheless, from A's perspective, continuing to fight is costly, so A will not always try to get the best bargain possible due to the costs associated with forcing B to update its beliefs.

When A's and B's beliefs differ drastically, A's expected gains from improving its future bargaining position outweigh the cost of fighting an additional period before reaching a deal. In contrast, when A's and B's beliefs are "close," then A's expected gains in future negotiations from bringing B's beliefs somewhat closer do not offset the additional cost of fighting, and A would therefore prefer to strike the bargain today.

In equilibrium, A only makes offers equal to $C_B(X_t, t)$. Yet, if A believes B is too optimistic about its chances of success, then A postpones making realistic offers and continues fighting. Fighting reveals more information about p and so causes the beliefs of the two parties to converge. As these beliefs converge, the additional amount A expects to extract from B in future bargains is insufficient to offset the cost of an additional battle. At such a point, A offers B the deal $b_t = C_B(X_t, t)$, which B accepts, ending the game.

CONVERGENCE OF BELIEFS

As the game progresses, with neither side achieving decisive victory, the two sides' beliefs about the mean value of p converge, and the variance associated with their estimates declines. In our numerical example, for example, initially A and B both believe their expected probability of winning the next battle is 2/3. However, over time, if the war does not end, then their beliefs grow closer together. In Figure 2, we follow a possible path of the game and plot A's and B's beliefs for each nation.

The solid line represents A's beliefs,

$$f\left(p; \alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t - X_t + X_0}{2}\right),$$

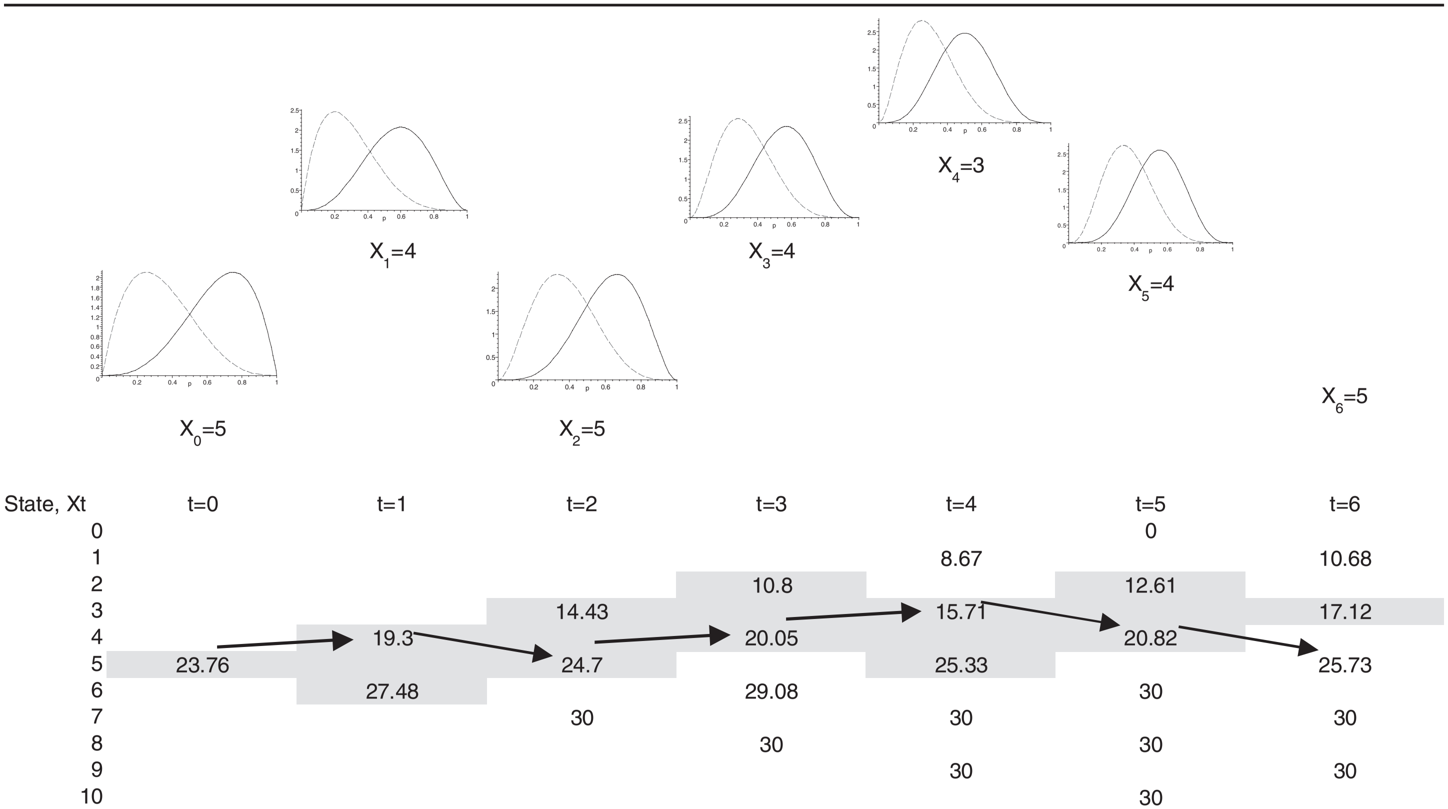

Figure 2: How Beliefs about the Value of *p* Change as the Game Progresses.

NOTE: The full and dotted lines represent A's and B's beliefs, respectively, for the sample path shown.

and the dotted line represents B's beliefs,

$$f\left(p;\alpha_{B0}+\frac{t+X_t-X_0}{2},\ \beta_{B0}+\frac{t-X_t+X_0}{2}\right).$$

As the figure shows, the longer the war lasts, the less the divergence of opinion about the true value of p. Lemma 1 is a technical result that states that A's and B's beliefs converge.

Lemma 1: For all $\varepsilon > 0$, there exists a $\tilde{t}(\varepsilon)$ such that for all $t \geq \tilde{t}(\varepsilon)$, and for states

$$X_t \in \{1, \ldots, N-1\}, |C_B(X_t,t)-\widehat{C_B^A(X_t,t)}| < \varepsilon.$$

Proof. This is a standard limiting argument. Take ε fixed.

$$C_B(X_t,t)-\widehat{C_B^A(X_t,t)}=\int_0^1 (V\lambda(p,X_t)-KE[R|p,X_t])[f(p;\alpha_{Bt},\beta_{Bt})-f(p;\alpha_{AT},\beta_{At})]dp.$$

Let $H(X_t) = \max_p(|V\lambda(p, X_t) - KE[R|p, X_t]|)$, and let

$$g(X_t,t)=\left|f\left(p;\alpha_{B0}+\frac{t+X_t-X_0}{2},\ \beta_{B0}+\frac{t+X_t-X_0}{2}\right)\right.$$
$$\left.-f\left(p;\alpha_{A0}+\frac{t+X_t-X_0}{2},\ \beta_{A0}+\frac{t+X_t-X_0}{2}\right)\right|.$$

Therefore, $|C_B(X_t,t)-\widehat{C_B^A(X_t,t)}| \leq H(X_t)\int_0^1 |(f(p;\alpha_{Bt},\beta_{Bt})-f(p;\alpha_{At},\beta_{At}))|dp$. Yet because $f(p;\alpha_{Bt},\beta_{Bt})$ and $f(p;\alpha_{At},\beta_{At})$ converge in distribution, there exists some $\hat{\tilde{t}}(\varepsilon,X_t)$, such that $\int_0^1 g(X_t,t)dp \leq \varepsilon / H(X_t)$ for state X_t. Let $\tilde{t}(\varepsilon) = \max \hat{\tilde{t}}(\varepsilon, X_t)$. Note as $\tilde{t}(\varepsilon)$ becomes large, beliefs converge to $p = 1/2$.

That is, if the game continues sufficiently long without either side attaining a decisive victory, then A's and B's beliefs will converge. Our next result exploits this property. Once the two sides' beliefs have converged sufficiently, A's leaders will prefer to make a deal today rather than delay another period, which is costly. Going one extra period does not offer the prospects of improving A's payoff by more than K, so A makes the deal today.

Proposition 2: In all $t > \bar{t} = \tilde{t}(K)$, A offers $b_t = \max\{0, C_B(X_t, t)\}$, which B accepts.
Proof. If $C_B(X_t, t) < 0$, then A attains the whole prize. We consider only the more difficult case of $C_B(X_t, t) \geq 0$.

Suppose A offers $b_t = C_B(X_t, t)$ in every period $t \geq \bar{t}$; then, B's expected value of rejecting A's offer is $C_B(X_t, t)$. Hence, B accepts if $b_t \geq C_B(X_t, t)$ and rejects if $b_t < C_B(X_t, t)$. In common with many similar games, there is not a subgame-perfect equilibrium in which B accepts $b_t > C_B(X_t, t)$

and rejects $b_t \leq C_B(X_t, t)$ because, for any offer $b > C_B(X_t, t)$, nation A can always find an offer b' such that $b > b' > C_B(X_t, t)$. Given A's offer $b_t = C_B(X_t, t)$, B's strategy is a best response.

We now consider A's strategy. Because B accepts all deals $b_t \geq C_B(X_t, t)$, A never offers any deal greater than $C_B(X_t, t)$. If A offers $b_t = C_B(X_t, t)$, then A's payoff is $M(X_t, t)$. If A offers $b_t < C_B(X_t, t)$, then A's expected payoff is $-K + E_A[p|X_t, t]M(X_t + 1, t + 1) + (1 - E[p|X_t, t])M(X_t - 1, t + 1)$, where

$$E_A[p|X_t, t] = \int_0^1 pf\left(p; \alpha_{A0} + \frac{t + X_t - X_0}{2}, \beta_{A0} + \frac{t + X_t - X_0}{2}\right)dt$$
$$= \frac{\alpha_{A0} + \frac{t + X_t - X_0}{2}}{\alpha_{A0} + \frac{t + X_t - X_0}{2} + \beta_{A0} + \frac{t + X_t - X_0}{2}}.$$

But this payoff is less than

$$J = -K + V - E_A[p|X_t, t]\widehat{C_B^A(X_t + 1, t + 1)}$$
$$-(1 - E[p]X_t, t])\widehat{C_B^A(X_t - 1, t + 1)} = -K + V - \widehat{C_B^A(X_t, t)}.$$

Intuitively, J is the expected amount A could extract from B tomorrow by winning a single additional battle, thereby causing B's beliefs to converge with A's beliefs, less the cost of the additional battle. But, by lemma 1, for $t \geq \bar{t}$, $M(X_t, t) > J$. That is, even if an additional battle altered B's beliefs, bringing them fully in line with A's, the additional improvement in A's bargaining leverage is worth less than the cost of an addition battle. Hence, A offers $b_t = \max\{0, C_B(X_t, t)\}$.

Corollary: The smaller K, the larger $\bar{t}$.

Proposition 2 ensures that the game ends within some finite time either because one side decisively defeats the other, or, at time $\bar{t}$, beliefs have sufficiently converged that a negotiated settlement is inevitable. Given that the game is finite, we can solve the game by simple backward induction. We characterize A's continuation value, $\theta(X_t, t)$, from being in state X_t at time t. These continuation values are the additional expected payoff from playing the game starting at state X_t at time t. In the absorbing state 0, A has already decisively lost the war, so A's continuation value is $\theta(0, t) = 0$ for all t. Similarly, if A has already decisively won the war, $\theta(N, t) = V$ for all t. Proposition 2 also allows us to assign continuation values for time period $t \geq \bar{t}$ because we know that a deal will be reached without further conflict. Hence, $\theta(X_t, t) = M(X_t, t) = \max\{0, V - C_B(X_t, t)\}$ for all $t \geq \bar{t}$. Proposition 3 characterizes the continuation values in all time periods and states via backward induction.

Proposition 3: For $t \geq \bar{t}$, $\theta(X_t, t) = M(X_t, t)$. For $t < \bar{t}$, A's continuation values are defined by backward induction: $\theta(X_t, t) = \max\{M(X_t, t), -K + E_A[p|X_t, t]\theta(X_t + 1, t + 1) + (1 - E_A[p|X_t, t])\theta(X_t - 1, t + 1)\}$, where

A's continuation values for each state.
In the shaded states A prefers to continue the conflict rather than meet B's minimum bargaining demands.

State, X_t	t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7	t = 8	t = 9	t = 10
0						0		0		0	
1					8.67		10.68		12.36		13.77
2				10.8		12.61		14.16		15.65	
3			14.43		15.71		17.12		18.42		19.6
4		19.3		20.05		20.82		22.09		23.19	
5	23.76		24.7		25.33		25.73		27.1		28.2
6		27.48		29.08		30		30		30	
7			30		30		30		30		30
8				30		30		30		30	
9					30		30		30		30
10						30		30		30	

A's rewards from a negotiated settlement in each state.

State, X_t	t = 0	t = 1	t = 2	t = 3	t = 4	t = 5	t = 6	t = 7	t = 8	t = 9	t = 10
0						0		0		0	
1					8.67		10.68		12.36		13.77
2				9.71		12.12		14.06		15.65	
3			11.33		14.22		16.44		18.18		19.6
4		13.87		17.26		19.69		21.53		22.98	
5	17.78		21.48		23.94		25.73		27.1		28.2
6		26.92		29.08		30		30		30	
7			30		30		30		30		30
8				30		30		30		30	
9					30		30		30		30
10						30		30		30	

Figure 3: A's Continuation Values ($\theta(X_t, t)$) and the Best Bargain A Can Reach ($M(X_t, t)$)
NOTE: In the shaded states, A prefers to continue the conflict rather than meet B's minimum bargaining demands.

$$E_A[p|X_t,t] = \int_0^1 pf\left(p;\, \alpha_{A0} + \frac{t+X_t-X_0}{2}, \beta_{A0} + \frac{t+X_t-X_0}{2}\right)dt$$
$$= \frac{\alpha_{A0} + \frac{t+X_t-X_0}{2}}{\alpha_{A0} + \frac{t+X_t-X_0}{2} + \beta_{A0} + \frac{t+X_t-X_0}{2}}.$$

Proof. By proposition 2, we know that for $t \geq \bar{t}$, the game ends with the negotiated settlement $b_t = \max\{0, C_B(X_t, t)\}$. Hence, $\theta(X_t, t) = M(X_t, t)$.

For $t < \bar{t}$, consider the following strategy: B accepts b_t if $b_t \geq C_B(X_t, t)$ and rejects otherwise. A offers $b_t \in \{C_B(X_t, t), 0\}$. We now show that these are best responses. Given $b_t \in \{C_B(X_t, t), 0\}$, B never receives an offer greater than $\max\{C_B(X_t, t), 0\}$ in any future period. Hence, accepting b_t if $b_t \geq C_B(X_t, t)$ is a best response in every subgame. We now consider A's optimal strategy. Given B's acceptance strategy, any offer $b_t \geq C_B(X_t, t)$ is accepted. Offering $b' > \max\{C_B(X_t, t), 0\}$ is dominated by offering b'', such that $b' > b'' > \max\{C_B(X_t, t), 0\}$. Therefore, any offer strategy by A must be either $b_t = \max\{C_B(X_t, t), 0\}$ or an offer that is rejected.

Consider $t = \bar{t} - 1$. If A offers $b_t = \max\{0, C_B(X_t, t)\}$, then A's payoff is $M(X_t, \bar{t} - 1)$. We have already shown that A never makes a larger offer, so the only other case to consider is an offer that will be rejected, $b_t < C_B(X_t, t)$. In this case, A's payoff is

$$-K + \Pr(X_{t+1} = X_t + 1)\,\theta\,(X_t + 1, t + 1) + \Pr(X_{t+1} = X_t - 1)\,\theta\,(X_t - 1, t + 1) =$$

$$-K + E_A[\,p|X_t, t]\,\theta\,(X_t + 1, t + 1) + (1 - E_A[\,p|X_t, t])\,\theta\,(X_t - 1, t + 1).$$

A's continuation value is the larger of the expected value from an additional battle or a negotiated settlement today.

Next consider $t = \bar{t} - 2$. By the same logic as above,

$$\theta(X_t, t) = \max\{M(X_t, t), -K + E_A[p|X_t, t]\,\theta\,(X_t + 1, t + 1) + (1 - E_A[p|X_t, t])\,\theta\,(X_t - 1, t + 1)\}.$$

Repeating this argument for $t = \bar{t} - 3, \bar{t} - 4, \ldots, 0$ defines the continuation values for each nation and time period. Figure 3 demonstrates the logic of the backward induction argument made in proposition 3.

The top portion of Figure 3 shows the continuation values $\theta(X_t, t)$, and the lower portion, $M(X_t, t)$, shows the largest share of the prize that A could extract from B via negotiations; at some point in the war, these two values intersect. The exact point depends on the outcomes of the actual battles. In earlier periods, the values for $M(X_t, t)$ and $\theta(X_t, t)$ often differ, with $\theta(X_t, t)$, the value of continuing to fight, being larger. In the example in Figure 3, at $t = 0$, A values fighting 23.8 and settling 17.8, so A fights. Depending on the outcome of the first two battles, at $t = 2$, A either prefers to settle (if A loses the first two battles) or continues to fight (if A either wins one or two of the first two battles). However, as the game progresses, these numbers converge. By time period $t = 10$, $M(X_t, t) = \theta(X_t, t)$ for all X_t. This means that if a decisive victor has not emerged by $t = 10$, and a negotiated settlement has not already been reached, then an agreement will be reached in $t = 10$. At $t = 9$, $M(X_t, t) = \theta(X_t, t)$ for states 2, 6, and 8 but not for state $X_t = 4$, where $\theta(4, 9) = 23.19 > M(4, 9) = 22.98$. Referring to Figure 2, we see that at $X_9 = 4$, B believes the value of p to be $(\alpha_{B0} + 4)/(\alpha_{B0} + 4 + \beta_{B0} + 5) = 0.4$. After losing four battles and winning five, B believes that it has a 0.8 probability of being the eventual winner if the war is fought to its eventual conclusion, which B believes will take, on average, about additional 17 periods. Hence, B's expected value for fighting the war is about 7. A can offer B a 7-unit share of the prize, end the war, and keep the remaining 23-odd units of the prize for A. Yet, A prefers to fight because, from A's perspective, the continuation value is greater than B's reservation value.

A's leaders estimate their probability of winning the next battle as $E_A\,[p \mid X_9 = 4] = (\alpha_{A0} + 4)/(\alpha_{A0} + 4 + \beta_{A0} + 5) = 0.53$. As we have already seen, in period $t = 10$, an agreement will be reached. If A wins the next battle, its leaders can extract $\theta(X_{10} = 5, 10) = M(X_{10} = 5, 10) = 28.2$ from B. Alternatively, if A loses an additional battle, tomorrow it can extract only $\theta(X_{10} = 3, 10) = M(X_{10}\,3, 10) = 19.60$. Hence, fighting one additional battle before reaching an agreement with B tomorrow is worth $(0.53 * 28.2) + (0.47 * 19.60) - K = 23.2$. Because of this, at $X_9 = 4$, A believes that the expected bargaining advantage from an additional battle offsets the additional cost, and so A fights.

Repeating this backward induction argument provides the continuation values shown in Figure 3. The shaded states in Figure 3 represent those states in which A prefers an additional battle rather than to reach an agreement today.

WHY PRIORS DIFFER IN THE FIRST PLACE

Our conflict model shows how differences in prior beliefs can lead to war and, in turn, how fighting battles resolves these differences, at least partially. To complete our picture of international relations, we need to explain how beliefs, once converged, can again diverge. How can two rational actors, observing the same information, reach different conclusions about their relative strength? Psychological and bureaucratic theories provide some insights.

Psychologists recognize that a nation's leaders tend to develop theories of warfare whereby the nation's success relies heavily on the factors perceived to be the ones in which it holds the greatest advantage over its opponent. As Blainey (1973, 40) explains,

> In England the prediction that the war of 1914 would be short was based heavily on the economic arguments. England was the leading financial power: accordingly, if economic collapse was to come early in the war, it would hit England's enemies first and so lead to their surrender. In contrast, German leaders predicted that the war would be short because of the decisiveness of modern military technology: in that field, Germany was the recognized master and so could expect victory. Expectations of the outcome of the war had a strong, subjective, inarticulate streak.

Such explanations encompass the common phenomenon that the vast majority of people regard themselves as good drivers. Although initially it appears impossible that most are better than average, when the speedy equate good driving with speed, the cautious regard good driving as careful driving, and the skillful liken keen handling to good driving, then it is easy to see that, via the use of differing perceived appropriate standards, each plausibly regards themselves as "above average."

Bureaucratic politics can create similar biases (Allison 1999). Because the agencies reporting on military preparedness have their own agendas, their reporting tends to be biased in a manner designed to support their goals. For instance, before and during the initial phases of the First World War, the French military adopted an offensive military doctrine grounded in *élan*, or an aggressive martial spirit (Snyder 1984). Such a doctrine provided the military with great freedom to structure the military to suit its goals. Although these goals presumably included national defense, the military might have other objectives, commonly including increasing its share of the budgetary pie (Allison 1999; Kier 1997; Posen 1984). Thus, even without any misperception on the part of the military, it may present biased information to civilian policy makers.

Although psychological and bureaucratic explanations provide compelling explanations for divergent beliefs, they are also outside the theoretical framework of our model. Although either of these mechanisms is sufficient to explain why nations choose to disagree, we prefer to explain the origin of divergent beliefs within the context of the model. That is, we want to explain why fully rational actors, who update their beliefs according to the Bayes rule, can take away different lessons from the same event, as was the case in the events prior to the Seven Weeks War.

Before the Seven Weeks War, both Austria and Prussia (among others) observed the devastating effect of firepower during the American Civil War. Although both sides saw the same evidence, they drew different conclusions (Luvaas 1959). The Prussians observed the devastating power of massed fire. So did the Austrians. Where they differed was in how they thought the development of the needle gun affected their relative strength. All of the European observers fit the information they gleaned from the American Civil War into their preexisting models of warfare. Given the lessons of the American Civil War, the Prussians felt emboldened by their widespread deployment of the needle gun. Although the Austrians also saw the power of concentrated fire, they felt that it flowed from highly disciplined and cohesive units, rather than the technological innovations of the needle gun. They believed that although the needle gun gave the individual soldier greater firepower, it undermined unit cohesion because soldiers would rapidly discharge all their ammunition and then retire to the rear (Wawro 1997). Although both arguments have intrinsic logic, the evidence of the Seven Weeks War suggests that the improvement in firepower overwhelmed any loss in unit cohesion.

FORMING BELIEFS ABOUT *p*

We use the Austria and Prussia example as the departure point for our investigation. Just as political scientists and economists disagree over the relative value of different methodological approaches, nations can disagree over the origins of military power. For example, the Correlates of War (COW) project developed a composite index of national strength (Singer, Bremer, and Stuckey 1972). Although a useful background, there is considerable controversy about whether the five factors (and their relative weighting) in the index really reflect the true strength of nations. What is more, what constitutes strength in one context might be a liability in another. For instance, Stam (1996) demonstrates that military strategy is the most important determinant of war outcomes, but the effectiveness of different military strategies and technologies is heavily dependent on the terrain and location of the war. Hence, rather than thinking entirely in terms of p, which was sufficient in the above model of belief convergence, to explain differential learning and subsequent divergence of beliefs, we need to think about the components that go into p.

Let Z represent a vector of all the components that could conceivably affect military outcomes. In reality, this is a very long list of interactive factors, but for current purposes, we shall assume that it is composed of only three: $Z = (Z_{red}, Z_{yellow}, Z_{blue})'$. The extent to which these factors affect outcomes depends on a row vector of parameters, θ, and a link function, $\Phi(\theta Z)$. Probit, a common econometric model for dealing with binary data, fits within this framework and is the setup we use for our example. Here, $\Phi(\theta Z)$ is the cumulative distribution of a standard normal random variable evaluated at θZ. Hence, $p = \left(1/\sqrt{2\pi}\right) \int_{-\infty}^{\theta Z} e^{-x^2/2}\, dx$.

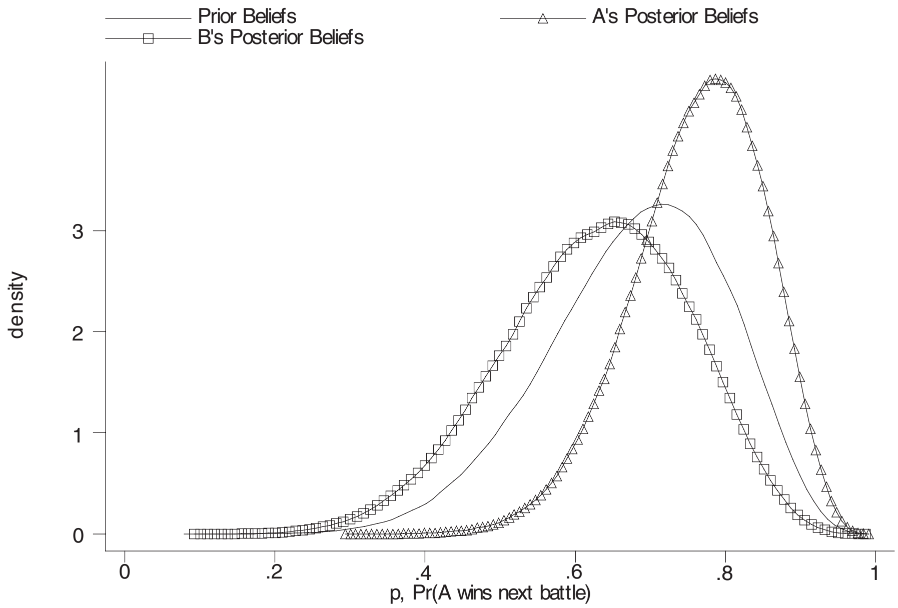

Figure 4: Divergence in A's and B's Beliefs following Five Victories by C over D

Just as in other aspects of life, nations with noncommon priors hold differing theoretical perspectives of warfare. Suppose, for example, that nation A subscribes to the orange theoretical perspective, thinking that red and yellow factors are important in determining battle outcomes but blue factors much less so. In contrast, suppose nation B is a green theorist, believing yellow and blue factors are the most critical. To construct a concrete example, from A's orange perspective, A's beliefs are that θ is normally distributed with mean (1, 1, 0)′ and a variance of 1 for the first two components and zero for the last:

$$\theta \sim N_3\left(\begin{pmatrix}1\\1\\0\end{pmatrix}, \begin{pmatrix}1 & 0 & 0\\0 & 1 & 0\\0 & 0 & 0\end{pmatrix}\right).$$

In contrast, from B's green perspective,

$$\theta \sim N_3\left(\begin{pmatrix}0\\1\\1\end{pmatrix}, \begin{pmatrix}0 & 0 & 0\\0 & 1 & 0\\0 & 0 & 1\end{pmatrix}\right).$$

Suppose that, in the context of a recent war between A and B, the relevant factors are $Z_{AB} = \left(Z_{red} = \frac{1}{4}, Z_{yellow} = \frac{1}{4}, Z_{blue} = \frac{1}{4}\right)'$. Perhaps having recently fought a war, nations A and B have no disagreement about their relative strength: both agree that $E[p] = .68$. Their identical prior beliefs over the distribution of p appear in Figure 4. With no disagreement over the relative value of p, A and B can agree on a bargain.

Suppose A and B witness a war or important military event involving C and D. It is important to note that the event need not be as dramatic as a war. Even a failed missile test might create differential learning. For instance, one side might believe that a faulty guidance system caused the failure, whereas the other side might discount the failure and instead draw a positive impression due to the missile's speed or payload capacity. C and D are unlikely to have exactly the same distribution of military factors as A and B, and so the implications for A and B of a victory by C must reflect the two states' different capabilities. Suppose $Z_{CD} = \left(\frac{1}{2}, \frac{1}{4}, -\frac{1}{2}\right)'$ and that A and B both witness five straight victories for nation C.

If $g_A(\theta)$ is A's probability density function over θ, then upon seeing C's five victories (data Y), A's beliefs shift to $g_A(\theta|Y)$, where $g_A(\theta|Y) \propto g_A(\theta)\Phi(\theta Z_{CD})$ by the Bayes rule.[7] Because of the information revealed during the war between C and D, A's and B's beliefs over θ shift but not in identical ways. In particular, rather than believing that the expected impact on A's chances against B of red and yellow is 1.0, A now believes that it is 1.65 and 1.33. A's beliefs about the effects of blue remain unchanged. Similarly, for B, $E_B[\theta|Y] = (0, 1.637, -.274)$. Although it is interesting that nations' shift their beliefs about the relative influence of different factors on battle outcomes, our primary concern is how this alters their prior agreement about p.

Figure 4 shows A's and B's posterior beliefs about p, the probability that A would win the next battle. Having seen five victories for C, A thinks the expected value of p is .76, whereas B thinks the expected value of p has fallen to .63. In short, both sides now believe that they are stronger, as well as stronger relative to their potential opponent, having seen the same information—C's five consecutive successes. Skeptics might think these differences occur solely because of the radical differences that A and B place on red and blue. Suppose, instead, that $Z_{AB} = \left(0, \frac{1}{4}, 0\right)'$, so that differences in red and blue are of little consequence, and p is driven only by yellow, a factor that both agree is important. Having observed five victories for C, A's and B's expectations over p are .63 and .66, respectively. Although the effects are smaller, A's and B's beliefs about p still diverge.

Of course, the extent to which observing five victories for C causes A's and B's beliefs to diverge depends on the strength of their beliefs in the first place. The tighter their prior beliefs about p are, by which we mean the smaller the variance in each side's beliefs about p, the less will be the impact of new information. All else equal, in terms of causing A's and B's beliefs to diverge, new information has less impact following a long war composed of many battles than a short (and, at the time, apparently decisive) conflict that provided little opportunity for learning.

7. These posterior densities are often analytically intractable. We used a Markov chain Monte Carlo simulation method to calculate them (Albert and Chib 1993).

Although nations come to agree on their relative strength through conflict, convergence in beliefs about p does not necessarily imply convergence in beliefs about θ, the factors that determine p. For the sake of the above example, we supposed that nations disagree on the theoretical nature of future warfare. Such an assumption is consistent with empirical reality and the economic literature on common priors noted earlier. Given different theoretical perspectives on the nature of war, when nations observe data that, although related, do not come from a direct contest between A and B, they can draw different inferences about their own likely future wartime performance (Luvaas 1959). This differential learning permits a divergence of beliefs about nations' relative strength to develop.

TECHNOLOGICAL CHANGE

A skeptic might argue from the perspective of the private-information model of disagreement that although the orange and green theoretical perspectives led initially to different conclusions about p, eventually enough information would accumulate so that the nations would agree as to the "correct" theory. One might conceptualize this as Bayesian model selection (Poirier 1995, chap. 7). If the orange theory were consistently to outperform the green theory, then nations would slowly shift toward believing the predictions of the orange theory. We might then logically propose, given the unfortunately vast history of warfare, that nations should agree by now on the best model of war. Unfortunately, such a conclusion requires the absence of technological innovation over time. As technology changes over time, so too must our theories of tactics and warfare. Theories of battle and war outcomes that were valid at one time will become obsolete due to subsequent changes in technology, thereby reopening room for rational disagreement about future outcomes.

For example, while English superiority in archery enabled them to win famous victories over the French at Crecy and Agincourt during the Hundred Years War, this information is useless in assessing the likelihood of an English victory against the French today. Technological changes shift the relative salience of different factors that influence battle outcomes and may even necessitate the inclusion of previously unknown or unobserved factors into nations' assessment of their relative strength. Such technological evolution prevents the consistent convergence on the best, or most accurate, theory of war.

Technological shifts mean that the impact of factors on the outcome of war changes over time: θ changes. Hence, progress toward information accumulation and its concomitant convergence of beliefs is limited. As θ changes, older information based on a θ that is now less relevant for assessing relative strength becomes obsolete. As with our arguments on differential learning, we demonstrate this formally. To model continually changing technology, we treat θ as a random walk. For instance, $\theta_{t+1} = \theta_t + \varepsilon_{t+1}$, where ε_{t+1} is some small change in the vector θ over a unit of time. Consistent with many standard models (e.g., Brownian motion), assume $\varepsilon_{t+T} \sim N(0, T\zeta^2)$; that is, shifts over unit time are normally distributed with mean zero and variance ζ^2. Even if a nation is certain that the value of θ at time t is θ_t, then T units of time later, its beliefs would be distributed normally with mean θ_t and variance $T\zeta^2$. As time progresses, a nation becomes less and less certain about θ, absent any additional information. Not only do nations have the problem of inferring θ, but they must also grapple with the fact that θ is a moving target.

Our story is now complete. We have explained that war occurs when nations disagree about their relative power, and we have shown how such disagreements can arise between completely rational actors.

WAR TODAY, WAR TOMORROW?

One insight that flows from our model builds on the obvious logic that bargains are not attainable until the two nations' beliefs converge sufficiently so that they would prefer to reach an agreement today rather than fight another battle. The puzzle is, how far must the two sides' beliefs converge to bring peace in a war fought today, and how does this necessary degree of convergence, in turn, affect the likelihood of future conflicts occurring? How closely the two sides' beliefs need to converge depends on several factors, not the least of which is how far apart their beliefs are in the first place. Perhaps most important, the higher the cost of conflict, K, the less willing nations will be to fight one extra battle. Similarly, as nations approach an absorbing state, the quicker they are to reach an agreement.

In Figure 2, we showed how the two sides' previously divergent beliefs converge as a war progresses. Holding factors such as the cost of war and the gains from victory constant, the longer the war lasts, the more the two sides' beliefs will shift in two important ways. First, ceteris paribus, the longer the war, the greater the convergence of the two sides' beliefs about the mean war outcome and, second, the smaller the variance in their respective estimates of p. The greater the convergence, the more the two sides agree. The smaller the variance, the more certain each side is about its estimate of future battle outcomes. Combining the two factors, the greater the degree of agreement and the more certain the two sides are about that agreement, the larger the stochastic shock will have to be for the two sides to agree that war is preferable to peace.

The shaded area of Figure 2 represents states in which nations fight. Although it is obvious that not all wars last the same length of time, calculating the effect of the duration of war fought today on the likelihood of future war is not a trivial problem. For example, a war fought today might last only 2 periods, ending at $X_2 = 7$; alternatively, the war might last 6 periods, ending at $X_6 = 5$. Indeed, in our simple example, the war might last as long as 10 periods. Obviously, the longer the war, the more that A and B learn about p, one characteristic of the balance of power between A and B.[8]

8. This points out that natural language notions such as the balance of power are potentially problematic in that they are quite ambiguous. In the context of our model, the balance of power might refer to p,X or A's or B's beliefs about p,X. With regards to a conflict with its Arab neighbors, Israel has a high probability of winning the next battle (i.e., a large p). Yet, given its geography, if Israel loses even a few battles, it is decisively defeated (i.e., X_t is small).

Combining the learning that takes place within conflicts with how informational disagreements lead to war provides a link between how conflict today affects conflict tomorrow. First, we consider how information learned in an initial dispute affects the likelihood of conflict in a second dispute. Second, we examine how the length of the initial conflict affects the robustness of the settlement and whether perturbations in beliefs following the settlement of a war make a reopening of the conflict more or less likely in the future.

We begin by supposing that the first war has ended with a negotiated settlement. This outcome reflects the two sides' agreement over the likely costs and outcomes of future battles. Stochastic shocks, or changes to θ, and differential learning awaken the possibility of renewed fighting. Some settlements will be more susceptible to future shocks than others. Following the end of a war, nations start to reassess what the likely outcome of reopening conflict might be. This anticipated war outcome might not be the same as the outcome reached in the first conflict, thereby providing incentives to renegotiate.

A change in the state variable may also provide strong incentives for a nation to make new demands. Although nations may strike a deal in state X_t, the distribution of stakes being contested often reverts to the status quo ante bellum, X_0. For instance, the U.S.-Iraq war in 1991 ended with a negotiated settlement reached after the United States and coalition forces had driven deep into Iraq (Summers 1992). Because the key terms of the agreement did not include territorial concessions (at least with respect to Iraq sovereignty), the U.S. forces withdrew. Iraq was much closer to decisive defeat in February 1991 than it had been just a few months later. With decisive defeat more distant, the attractiveness of the deal reached to end the fighting diminished significantly for Saddam Hussein, who subsequently began to renege on many of the terms of the bargain that ended the fighting in 1991.

Holding other factors constant, the closer nations' beliefs about p and the more accurately that the distribution of forts resembles the state variable at the end of the conflict, the more stable and longer lasting any agreement should be (Box-Steffensmeier, Reiter, and Zorn 2003). In contrast, the greater the extent to which A and B disagree about the value of p, and hence who will win the second war, the harder it is for A and B to negotiate a peaceful settlement to the demands triggered by the shock. In contrast, if A's and B's beliefs are identical, they can immediately strike a new deal. Hence, the length and intensity of the initial conflict influence the likely course of subsequent disputes because they systematically affect the variance, or relative certainty, of the two sides' subjective beliefs about p. The greater the number of battles and the more intense they are in the initial war, the greater the convergence of beliefs and the smaller the variance in the two sides' estimates of future battle outcomes. In contrast, the shorter the initial war and the fewer the costs each side bears, the less each side will have learned, and as a result, the subsequent convergence of the two sides' beliefs will be less than in longer, more costly wars.

Although both sides update their beliefs during conflict, as a practical matter, conflict often has a much greater impact on one side's beliefs than on the other side's beliefs. For instance, during the recent war between the North Atlantic Treaty Organization (NATO) and Serbia, the long string of NATO bombing sorties without substantial losses probably shifted Serb beliefs to a far greater extent than those of NATO (Holbrooke 1998).

Disputes or demands for renegotiation that follow lengthy wars are more easily resolved without renewed use of force (Fortna 2004a, 2004b). Alternatively, other factors being equal—including the degree of disagreement about p prior to the first war—should war break out again, because the shock to the actors' beliefs about p was sufficiently large, the duration of the subsequent war is likely to be relatively short because much of the learning required for an agreement already took place in the previous war. During war, nations learn about p as they see the outcome of battles. When the two sides stop fighting, the nations end this common learning process. Yet, this does not imply that nations stop learning—they simply stop learning the same things from the same events. Nations continue to learn and update their beliefs about p from other sources. As a result, over time, nations' beliefs can drift away from the beliefs they held at the end of the war.

Again, consider Iraqi beliefs about U.S. capabilities, this time in 1990. Iraqi leaders believed the United States to be incapable of sustaining a large number of casualties while continuing a war effort. They also believed that the United States would be incapable of destroying the Iraqi military before it would be able to inflict substantial casualties on U.S.-led forces. The source of these beliefs traces back, in part, to the U.S. withdrawal from Vietnam and American civilian reaction to the accumulating casualties during the late 1960s and early 1970s (Gartner 1997; Gartner and Segura 1998). Although the Iraqi leaders' beliefs may or may not have accurately characterized U.S. attitudes in 1975, the Iraqis (as well as many of the pessimists in the U.S. Congress) failed to grasp the revolution in U.S. military capabilities that took place during the Reagan administration (Summers 1992).

Despite U.S. victories in every battle during the 1991 Gulf War, war between the United States and Iraq again became increasingly likely in late 2002, early 2003. This highlights a potential downside to the so-called Powell/Weinberger doctrine, which dictates the use of force only to achieve short and decisive victories accomplished with overwhelming forces, as compared to the just war doctrine, which advocates a more proportionate application of force. Although the rapidity of the U.S. military victory in 1991 fulfilled its immediate goals, it provided little opportunity for learning. For instance, it provided no evidence about how American civilians would react to a potentially higher level of causalities that would occur in battles fought in the close quarters of urban areas during a potential insurgency. Because the 1991 Gulf War did not reveal this information, and given the reversion of the distribution of forts close to the status quo ante bellum (the United States withdrew the majority of its forces from Iraq and the region), Saddam Hussein was willing to renege on the deal struck during the Gulf War almost immediately.

As we discussed above, nations might draw different lessons even from the same events. Therefore, over time, beliefs likely diverge, just as the initial priors in the game diverged. When an initial war is short, little information is revealed, which leads to the two sides' beliefs remaining relatively more diffuse than if the war had lasted longer. As a result, relatively smaller shocks or perturbations can more readily shift the beliefs of the nations out of line, making new demands more likely. As the two sides' beliefs begin to diverge, future disputes become harder to negotiate, and the nations eventually may feel sufficiently aggrieved by the outcome of the initial conflict that they reopen the initial dispute. However, when a conflict is long, with many battles waged, nations' posterior beliefs are much tighter. As such, compared to wars that have revealed little information, a much greater shock is required to swing beliefs sufficiently far apart to lead to a reopening of the conflict.

Suppose in our numerical example that a negotiated settlement was reached at $X_2 = 7$. A's and B's posterior beliefs are $\alpha_{At} = 6$ and $\beta_{At} = 2$ and $\alpha_{Bt} = 4$ and $\beta_{At} = 4$, which generate expected values for p (A's probability of victory) of $\frac{3}{4}$ and $\frac{1}{2}$, respectively. Note first that even though the two nations settled, they do not agree about the expected outcome of a future war. Following the war, should nation B receive a piece of information that it regards as equivalent to winning a battle, then its beliefs would shift to $\alpha_{Bt} = 4$ and $\beta_{Bt} = 5$, which drops B's expected value for p to 4/9 (0.44), a shift of 1/18 (0.055). Alternatively, suppose instead that the initial war had taken much longer to resolve and that it lasted 10 periods, finishing at $X_t = 5$. A's and B's expected values for p would be 9/16 (0.56) and 7/16 (0.44). The same additional perturbation in B's beliefs would give B a new expected value for p of 7/17 (0.41), which is a shift of less than half the magnitude of the shift in the first example. In a dispute following the long war, B would demand far less from A than it would following the shorter war.

This application of the bargaining model suggests several observable implications that are consistent with the existing empirical literature on international conflicts.

1. The higher the cost of conflict, the shorter the war, on average (Bennett and Stam 1996; Slantchev forthcoming).
2. The longer and more recent any previous conflict, the less likely any subsequent dispute is to escalate to conflict (Fortna 1998, 2004a, 2004b).
3. Should such a dispute escalate to conflict, the war is likely to be shorter in duration (Fortna 1998, 2004a, 2004b; Filson and Werner 2002, 2004; Slantchev forthcoming).

CONCLUSION

In this article, we presented a unified model of war initiation, duration, and recurrence. In particular, our model helps address the puzzle of how war today affects the likelihood of future demands for renegotiation of the postwar status quo and, hence, future war. By beginning with the following premise—that to understand the origins of war, we must understand the process of war—we reach a few novel conclusions. We find that wars characterized by very high-cost battles tend to be short, as the empirical work on war durations suggests (Bennett and Stam 1996; Slantchev forthcoming). These short but costly wars are also more susceptible to claims for renegotiation or unstable postwar equilibria when compared to longer wars, which reveal more information about the capabilities of the two sides, an argument consistent with the literature on dispute settlements (Werner 1998, 1999). Holding constant the myriad other factors associated with the onset of war (Bennett and Stam 2004), wars that are long and filled with relatively low-cost battles are more likely to lead to stable postwar settlements (Goertz and Regan 1997; Regan and Stam 2000; Box-Steffensmeier, Reiter, and Zorn 2003).

Our theory, if true, has significant policy relevance. Consider the potential effects of the so-called Powell Doctrine, one of whose tenets is that if the United States is to use force, it should be used overwhelmingly, thereby leading to wars of the shortest duration possible. This type of policy, although minimizing costs in the short run, also leads to situations in which the future is likely to be less stable than it might be otherwise.

Our model also leads us to a conclusion that differs from most bargaining approaches to international conflict. For example, Wagner (2000, 475) and others have claimed that when two nations are considering the use of force,

> the greater the per period costs of war and (therefore) the more likely it is to end quickly, the wider the range of status quo distributions that would satisfy both nations. Conversely, the less likely the war is to end quickly, the narrower the ranges of status quo distributions that are invulnerable to renegotiation by the use of military threat.

In other words, from this rationalist perspective, anticipated high-cost short wars lead to wider ranges of status quos that will remain stable once the fighting ends in the initial war. Anticipated long wars lead to narrower ranges of status quos that are comparatively unstable. In other words, in models built using the more standard setup, when nations expect wars to be costly but short, they should be willing to accept a broader range of negotiated solutions. When they think wars will be relatively low cost but of long duration, they should have a very narrow range of negotiated outcomes that would be acceptable (Fearon 1995; Wagner 2000; Filson and Werner 2002; Goemans 2000; Slantchev 2003).

By taking a different approach to modeling the nature of disagreement, we generate results that differ from the standard formal intuitions and that are more consistent with existing empirical results. Our model suggests, in contrast to much of the existing literature, that high-cost short wars will influence actors' beliefs less than wars that are low cost per period but longer. Wars that do little to force the convergence of the two sides' beliefs are far more likely to reopen than wars that lead to confident agreement about the two sides' relative capabilities.

Which story finds greater support in the empirical literature? Some of the existing evidence suggests that there is a powerful association between the nature of wars and the duration of the peace that follows (Fortna 1998, 2004a, 2004b; Werner 1999; Box-Steffensmeier, Reiter, and Zorn 2003). The existing quantitative studies provide results consistent with the theory presented here. For example, both Fortna (1998) and Werner (1999) show that as the cumulative costs associated strongly with the duration of a previous war rise, the duration of postwar peace increases. In contrast to Wagner's (2000) theoretical claim but consistent with the model presented here, Werner's empirical analysis demonstrates that the "settlements of wars with relatively low military costs are particularly vulnerable. Decreasing the value of the military cost variable by two standard deviations below its mean value, for instance, increases the hazard of another war by 55 percent" (p. 930). In a similar vein, Box-Steffensmeier, Reiter, and Zorn (2003) show that peace is more likely to break down following stalemated outcomes, in which the two sides continue to disagree about each other's capabilities, rather than nonstalemated outcomes, in which the postwar distribution of capabilities is clear. They also find that this effect diminishes with the passage of time during the subsequent periods of peace, consistent with our story of belief divergence. These empirical findings are quite robust and are consistent with the model presented in this study based on disagreement occurring with noncommon priors.

Our approach to disagreement is not a trivial stylistic shortcut. In addition to expanding our understanding of the nature of war, we propose innovations in modeling disagreement between actors. The traditional, private-information approach models disagreements by assuming that, initially, all players have the same beliefs. Differences in beliefs are created by players receiving private information about the state of the world. One effect of the Harsanyi (1967-1968, 1995) doctrine is that game theorists have been led to privilege the role of uncertainty and private information, which in turn leads to an almost myopic study of bluffing, reputation, cheap talk, audience costs, and the credibility of international commitments (Fearon 1995; Powell 2004a, 2004b; Schultz 2001). In the world in which our model resides, disagreement results not from either side bluffing but from the two sides clinging to different theories about the nature of war. Our approach forces one to think less about the strategic incentives to misrepresent private information and more about how and why people disagree, even when confronted with a mutually agreed-on set of empirical facts. The heterogeneous-beliefs approach relaxes the common-priors assumption and instead replaces it with the restriction that players have different initial beliefs and that these beliefs are common knowledge for all players. Such an approach emphasizes how conflict solves belief asymmetries by providing common information to both parties. By resolving disagreements and creating convergence in beliefs, conflict enables states to resolve their differences.

REFERENCES

Albert, James H., and Siddhartha Chib. 1993. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88 (422): 669-79.

Allison, Graham. 1999. *Essence of decision: Explaining the Cuban Missile Crisis.* Boston: Little, Brown.

Aumann, Robert J. 1976. Agreeing to disagree. *The Annals of Statistics* 4:1236-9.

Bennett, D. Scott, and Allan C. Stam III. 1996. The duration of interstate wars, 1816-1985. *American Political Science Review* 90:239-57.

———. 2004. *The behavioral origins of war.* Ann Arbor: University of Michigan Press.

Bernheim, Douglas. 1986. Axiomatic characterizations of rational choice in strategic environments. *Scandinavian Journal of Economics* 88:473-88.

Blainey, Geoffrey. 1973. *The causes of war.* New York: Free Press.

Boge, W., and T. Eisele. 1979. On the solution of Bayesian games. *International Journal of Game Theory* 8:193-215.

Box-Steffensmeier, Janet, Dan Reiter, and Christopher J. W. Zorn. 2003. Non-proportional hazards and event history analysis in international relations. *Journal of Conflict Resolution* 47:33-53.

Brandenburger, A., and E. Dekel. 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* 59:189-98.
Fearon, James D. 1995. Rationalist explanations for war. *International Organization* 49:379-414.
Filson, Darren, and Suzanne Werner. 2002. A bargaining model of war and peace. *American Journal of Political Science* 46:819-38.
———. 2004. Bargaining and fighting: The impact of regime type on war onset, duration, and outcomes. *American Journal of Political Science* 48 (2): 296-313.
Fortna, Virginia Page. 1998. A peace that lasts: Agreements and the durability of peace. Ph.D. diss., Harvard University.
———. 2004a. *Peace time: Cease-fire agreements and the durability of peace*. Princeton, NJ: Princeton University Press.
———. 2004b. Does peacekeeping keep peace? International intervention and the duration of peace after civil war. *International Studies Quarterly* 48 (2): 269-92.
Gartner, Scott S. 1997. *Strategic assessment in war*. New Haven, CT: Yale University Press.
Gartner, Scott Sigmund, and Gary Segura. 1998. War, casualties and public opinion. *Journal of Conflict Resolution* 42:278-300.
Goemans, H. E. 2000. *War and punishment: The causes of war termination and the First World War*. Princeton, NJ: Princeton University Press.
Goertz, Gary, and Patrick M. Regan. 1997. Conflict management in enduring rivalries. *International Interactions* 22 (4): 321-40.
Gul, Faruk. 1998. A comment on Aumann's Bayesian view. *Econometrica* 66 (4): 923-8.
Harsanyi, John C. 1967-1968. Games with incomplete information played by Bayesian players. *Management Science* 14:159-82, 320-34, 486-502.
———. 1995. Games with incomplete information. *American Economic Review* 85 (3): 291-303.
Holbrooke, Richard. 1998. *To end a war*. New York: Random House.
Kier, Elizabeth. 1997. *Imagining war*. Princeton, NJ: Princeton University Press.
Lewis, Jeffrey B., and Kenneth A. Schultz. 2003. Revealing preferences: Empirical estimation of a crisis bargaining game with incomplete information. *Political Analysis* 11:345-67.
Luvaas, Jay. 1959. *The military legacy of the Civil War*. Lawrence: University of Kansas Press.
Mertens, J. F., and S. Zamir. 1985. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14:1-29.
Morris, Stephen. 1995. The common prior assumption in economic theory. *Economics and Philosophy* 11:227-53.
Morrow, James D. 1997. Strategy, indeterminacy, and the politics of war. Paper presented at Opening the Black Box of War Conference, June, University of California, Davis.
Nyarko, Y. 1990. Bayesian rationality and learning without common priors. Working Paper 90-45, C. V. Starr Center for Applied Economics, New York University.
———. 1991. Most games violate the Harsanyi doctrine. Working Paper 91-39, C. V. Starr Center for Applied Economics, New York University.
———. 1997. Convergence in economic models with Bayesian hierarchies of beliefs. *Journal of Economic Theory* 74 (2): 266-96.
Poirier, Dale J. 1995. *Intermediate statistics and econometrics: A comparative approach*. Cambridge, MA: MIT Press.
Posen, Barry. 1984. *The sources of military doctrine: France, Britain and Germany between the World Wars*. Ithaca, NY: Cornell University Press.
Powell, Robert. 2004a. Bargaining and learning while fighting. *American Journal of Political Science* 48 (2): 344-61.
———. 2004b. The inefficient use of power: Costly conflict with complete information. *American Political Science Review* 98 (2): 231-41.
Regan, Patrick, and Allan C. Stam. 2000. In the nick of time: Conflict management, mediation timing and the duration of interstate disputes. *International Studies Quarterly* 44:239-60.
Schultz, Kenneth. 2001. *Democracy and coercive diplomacy*. New York: Cambridge University Press.

Singer, J. David, Stuart Bremer, and John Stuckey. 1972. Capability distribution, uncertainty, and war, 1820-1965. In *Peace, war and numbers*, edited by Bruce Russett. Beverly Hills, CA: Sage.

Slantchev, Branislav. 2002. The power to hurt: Costly conflict with completely informed states. *American Political Science Review* 47:1-11.

———. 2003. The principle of convergence in wartime negotiations. *American Political Science Review* 47:621-32.

———. Forthcoming. How initiators end their wars. *American Journal of Political Science* 48 (4).

Smith, Alastair. 1997. The nature of warfare: Power, bargaining, super weapons, elections and the democratic peace. Paper presented at the Western Regional Meeting of the International Studies Association, October, Davis, CA.

———. 1998. Fighting battles, winning wars. *Journal of Conflict Resolution* 42 (3): 301-20.

Smith, Alastair, and Allan C. Stam. 2003. Mediation and peacekeeping in a random walk model of war. *International Studies Perspectives* 5 (4): 115-36.

———. Forthcoming. When likely losers go to war. In *New direction for international relations*, edited by Alex Mintz and Bruce Russett. New York: Lexington Books.

Snyder, Jack L. 1984. *The ideology of the offensive: Military decision making and the disasters of 1914*. Ithaca, NY: Cornell University Press.

Stam, Allan C. 1996. *Win, lose, or draw: Domestic politics and the crucible of war*. Ann Arbor: University of Michigan Press.

Summers, Harry G. 1992. *A critical analysis of the Gulf War*. New York: Dell.

Varian, Hal R. 1989. Differences of opinion in financial markets. In *Financial risk: Theory, evidence, and implications*, edited by C. Stone. New York: Kluwer.

von Clausewitz, Carl. 1976. *On war*. Michael Howard and Peter Paret, eds. Princeton, NJ: Princeton University Press.

Wagner, R. Harrison. 2000. Bargaining and war. *American Journal of Political Science* 44:469-84.

Wawro, Geoffrey. 1997. *The Austro-Prussian War: Austria's war with Prussia and Italy in 1866*. Cambridge, UK: Cambridge University Press.

Werner, Suzanne. 1998. Negotiating the terms of settlement: War aims and bargaining leverage. *Journal of Conflict Resolution* 42:321-43.

———. 1999. The precarious nature of peace: Resolving the issues, enforcing the settlement, and renegotiating the terms. *American Journal of Political Science* 43 (3): 912-34.