Bargaining in the Shadow of War: When Is a Peaceful Resolution Most Likely?

Donald Wittman University of California

This article derives the optimal bargaining strategies of the belligerents when each side has private but incomplete information about the expected outcome of a war, should it take place. I show that the aggressor's demand curve can be below the defender's offer curve, that wars are possible even when both sides are jointly pessimistic, and that the relative cost of a war can radically alter the types of disputes that end in war. A simple diagram provides the intuition for most of the major propositions.

Which types of disagreements end in peace and which end in war? This is an old question with old answers. There have been endless theoretical arguments about whether preponderance of power or balance of power results in fewer wars. Powell (1999, 2002) demonstrated that much of the debate was a result of poorly defined models.[1] There is also a large body of contradictory empirical evidence supporting one or the other theory (again Powell is a useful resource). This article sheds new light on the question and provides an explanation for the contradictory results—the answer depends on the relative cost of war.

Here, war is viewed as a type of *common values* problem with private information. As an illustration of what is meant by common values, the outcome of a war depends on both defensive capabilities, which are known by the defender, and offensive capabilities, which are known by the aggressor. Surprisingly, except for Morrow (1989), the literature on war does not deal with the common values problem with two-sided private information. Instead, war is treated either as being a one-sided asymmetric information situation (for example, the aggressor knows the outcome of the war, but the defender does not) or as being a two-sided asymmetric information situation with private values (for example, each side knows its own cost of conducting the war but not the cost to the other side). So, this article also brings deeper understanding regarding the information structure and how the parties bargain in this more complicated information environment.

As will be seen, the model yields a number of counterintuitive results and generates a research agenda for future empirical research. The model can also serve as the basic structure for further theoretical refinements. Despite the formal presentation, the major ideas can be illustrated via a simple diagram. The outline of the article is as follows: In the next section, I model the bargaining situation, and, in the section after that, I provide an intuitive explanation for the major results. In later sections, I derive precise quantitative relationships, compare the results to the results of other models, consider various theoretical extensions, discuss how one might undertake empirical tests of the theory, and provide concluding remarks. In the appendix, I provide formal proofs of the major propositions.

The Framework of the Model

I start off with a discussion of the information framework.

War Is a Common Values Game Rather Than a Private Values Game

The outcome of a war depends on both the aggressor's offensive capabilities and the defender's defensive capabilities. And the same holds true for the cost of war to each

Donald Wittman is Professor of Economics, University of California, Santa Cruz, CA 95064 (wittman@ucsc.edu).

I thank seminar participants at Stanford and UCSD, the anonymous referees, and Bruce Bueno de Mesquita for very helpful comments.

[1] See also Wagner (1994).

American Journal of Political Science, Vol. 53, No. 3, July 2009, Pp. 588–602

 ISSN 0092-5853

side, which should be related in some way to the offensive and defensive capabilities of the warring sides. In general, one would expect that the defender would know more about its defensive capabilities than about the aggressor's offensive capabilities and that the aggressor would know more about its offensive capabilities than the defender's defensive capabilities.

All of this can be stated more formally. Let θ_A be the aggressor's privately known offensive capability. For example, the aggressor's private knowledge regarding the number of pilots it has trained is translated into θ_A, a measure of the aggressor's offensive capability. The defender does not observe θ_A but knows that it is drawn from a distribution function F^A. Continuing with our example, the defender does not know how many pilots the aggressor has, but has an idea that there are somewhere between 100 and 200 pilots (with each possibility in between being equally likely).

Symmetrically, the defender privately observes its own defensive capability, θ_D. As before, the other side has only imperfect knowledge. For example, the aggressor might not know how many rockets the defender has, but the aggressor does know the distribution of defensive power, F^D, from which θ_D is drawn.[2]

Here, the smaller θ_D is, *the greater the defender's capability*. Large values of θ_D and θ_A mean that the defender's defense is weak and the aggressor's offense is strong, which in turn means that the aggressor will do well if the conflict ends in war. We employ the following formula for the outcome of the war: $\bar{\theta} = [\theta_A + \theta_D]/2$. For now, it is convenient to assume that both θ_A and θ_D are bounded by 0 and 1. Because each side only has a partial picture of the outcome of the war, we can view θ_A and θ_D as being private signals about the *common value* $\bar{\theta}$. For example, if $\bar{\theta}$ is the percent of the land captured by the aggressor were there to be a war, then θ_A and θ_D are the private signals about the percent of the land that will be captured. Because we will be dealing in cost and benefits, we would treat $\bar{\theta}$ as the *value* of the land being captured; so, capturing half the land would be equivalent to capturing half the value. Giving up land is just a concrete example of the possible concessions. More generally, θ could be viewed as the *value of concessions* made by D.

[2]While some private information might be transmitted to the other side, the transmission is unlikely to be complete because each side might (1) want to exaggerate its strength or (2) be hesitant to provide information that could be useful for the enemy (e.g., that it has rockets hidden in mountain peaks). Hence, the private information that has been made public may shrink the boundaries of the distribution, but it is unlikely to shrink the distribution to a point. See Fearon (1995) for a discussion.

I use the phrase "common values" because the outcome of the war depends on the signals of both belligerents (the defensive strength of the defender and the offensive strength of the aggressor), as opposed to the usual characterization, where the outcome of the war depends on only one signal that is known by one side (the other side only knows that the signal is a random draw from a distribution). The phrase "common values" is typically associated with auctions, where failure to take into account the other bidders' expected observations results in the winner's curse—on winning, the bidder realizes that his or her estimate of the value of the object was the highest and therefore the winner overestimated the value of the item. To avoid the winner's curse, the person must take into account that all of the other bids and estimates will be below the winning bidder's estimate and discount appropriately. The equivalent of the winner's curse occurs in bargaining when a belligerent, say the aggressor, does not take into account the defender's expected observation, thereby misjudging the outcome of the war. Consequently, the aggressor demands too much when the aggressor has an above-average observation and demands too little when the aggressor has a below-average observation. As we will see below, when deciding how much to demand, the optimal strategy requires the aggressor to account for the defender's set of possible observations and offers.

We next consider the cost of the war to each side. It seems sensible to view the cost of the war to each side as being closely linked to how well the side will do. We assume that the cost to the defender increases and the cost to the aggressor decreases as $[\theta_A + \theta_D]/2$ increases. So, the cost to each side also depends on both signals rather than just one's own signal.

It is useful to compare this setup to the previous literature. Before doing so, I would like to point out that these authors have different objectives (undertaking an empirical study, getting into the black box of war, and so forth) and therefore have models that consider different issues from those I investigate. So, these models are more of an intersecting set than a subset of the model presented here. Furthermore, the work presented here (as well as some other published articles) can be seen as an extension of the work by Powell (1999), who employed the uniform distribution in analyzing war.

Powell (1999) models a situation where there is common knowledge regarding the outcome of the war, $\bar{\theta}$ (which he labels as p, the value of the land captured by the aggressor if there is a war) but where each side has private information about its own costs. So, there is two-sided incomplete information, but each side has complete information about what will happen to its own side in case there is war. This is a private values case because the one side's signal has no effect on the other side's payoff *if there is a war*.

A number of articles model one-sided asymmetric information. Filson and Werner (2002) and Reed (2003) assume that $\bar{\theta} = \theta_D$ is private information to the defender, but that the costs to each side are *common knowledge* (that is, the information is public). Slantchev (2003) assumes that one player knows θ and the other does not and that θ can take on only three values. As another example, Powell (2004) assumes that the defender is unsure of the aggressor's cost of fighting or probability of collapse, but the aggressor does have knowledge of these factors.

Finally, Morrow (1989) assumes that each side plus nature independently draw either a 0 or 1, with the outcome of the war being a win for the aggressor if at least two of the three draws equal 1 and a win for the defender, otherwise. He then restricts the negotiated settlements to three possible outcomes: 0, 1, or the status quo (which need not be close to the expected outcome of the war). So the information structure of all these models is more restrictive than the model considered here.

Bargaining Protocol

This article not only has a different information structure, but also a different bargaining protocol from the previous articles.[3] Reed (2003) assumes that the aggressor makes an offer and then the defender either rejects the offer and there is war or accepts the offer and there is a peaceful settlement to the conflict. Powell (1999) assumes that the defender makes the first offer and allows for the aggressor to make a counteroffer, but shows that no counteroffer will be made and instead the aggressor either goes to war or accepts the offer. Powell (2004) allows for a series of offers by the defender (the satisfied state in his terms) in a repeated game of negotiations with battles in between; while Filson and Werner (2002) have the aggressor make the offers over two time periods. Still another variation is found in Slantchev (2003), who allows for offer and counteroffer in a repeated stream of negotiations with battles in between. All of these can be viewed as screening by one side and/or signaling by the other. I consider a different protocol that enables me to elucidate the strategic thinking when there is a common values problem that I have described above.

Here, I present the following symmetric bargaining protocol. The aggressor submits its demand and the defender submits its offer to a neutral third party. If the offer is greater than or equal to the demand, then there is a settlement halfway between the two; otherwise, the aggressor goes to war.[4] The advantage of this characterization is that it reflects the symmetry of the knowledge structure. This is not a realistic characterization of the actual bargaining. But the protocol can be viewed as the reduced form of some more complicated but unspecified bargaining procedure when neither side has a bargaining advantage vis-à-vis the outcome of the war. For example, in the Rubinstein bargaining game with complete information, as the time between rounds goes to zero, the division of the surplus goes to one half, the same as we postulate for our game of incomplete information. One could use mechanism design to show that a symmetric division of the surplus implies the strategies that I derive here regardless of the protocol; that is, the bargaining strategies are "institution free" and do not depend on the protocol. The protocol is just an easy way of deriving the strategies.

The bargaining protocol follows in the path of the seminal work by Chatterjee and Samuelson (1983). They consider a buyer and a seller, each having a private value for the good drawn from a uniform distribution. If the demand by the seller is less than the offer by the buyer, then there is a trade at the halfway point; otherwise, the seller keeps the item and the buyer keeps her money. The Chatterjee-Samuelson model plays a central role in understanding bargaining in markets even though few, if any, actual buyer-seller negotiations employ the C-S protocol. One could imagine all kinds of complicated, possibly unsolvable, "realistic" dynamic protocols, where some information may or may not be revealed in earlier rounds. The reason for the influence of the C-S protocol is that it provides a sensible reduced form for these more complicated bargaining models at the same time it provides the actual solution to the particular protocol. I believe that the same argument holds for the protocol employed here.

Despite the similarities to Chatterjee and Samuelson and, to a lesser extent, Chatterjee (1981), this article differs in several fundamental ways. First, the truth can be revealed through a costly war, where the cost of war

[3]For surveys of bargaining in the shadow of war, see Reiter (2003) and Powell (2002).

[4]I confine the analysis to those cases where the *expected* payoff to the aggressor from going to war is positive (for discussion of when these conditions hold, see the section on extensions). One might ask why the third party does not (1) choose a settlement at the average of the demand and offer even if there is not an overlap or (2) figure out the true observations from the belligerents' demand and offer and then announce a settlement at the average of their observations. Unfortunately, that protocol is not incentive compatible as it provides perverse incentives, e.g., the aggressor demanding infinity.

to each side has both a fixed component and a variable component that depend on how poorly that side does in the war. Second, this article is concerned with common values where the true value depends on both signals; in contrast, the aforementioned articles deal with private values. Finally, I ask when is bargaining least likely to end in a settlement; Chatterjee and Samuelson and Chatterjee did not pursue this question.

Models where the uninformed side makes the offer and the informed side either accepts the offer or starts a war (e.g., Reed 2003) would not change their results if the bargaining protocol presented here (both sides present their offers simultaneously) were used instead. This is because the informed side knows what the uninformed side will offer (as it is based on expectations, known to all players). For example, if the informed side is the aggressor, the aggressor would demand the same amount as the defender will offer if the amount to be offered is greater than or equal to the net benefit of the aggressor going to war. Otherwise, the aggressor demands a higher amount. So this is the same result as would be the case if the offer were seen first. Since seeing the offer provides no new information to the aggressor, the simultaneous protocol will not change either player's bargaining strategy.

The Basic Structure of the Model

I now provide the basic structure of the model. There are two players, the aggressor country (A) and the defender country (D), and three stages. At the preliminary stage 0, the players have common knowledge about the structure of the game, including the payoff function and the distributions of signals. In particular, both players know that the worst possible outcome of the war for the aggressor is $L \geq 0$ and the best possible outcome is $U > L$. It is useful (but not necessary) to conceive of the war in terms of the amount of land captured by the aggressor. If there is a war, then the most land that the aggressor will capture is u with value U and the least amount of land that the aggressor will capture is l with value L. L will have to be sufficiently large so that, regardless of the signals, war is always a credible threat (in Powell's terms, we are looking only at cases where country A is dissatisfied). For mathematical convenience, I will initially assume that $L = 0$ and $U = 1$ (but in the section on extensions, I will shift L and U upward, and still later I will consider other variations, for example L = .8 and U = 1).

At stage 1, the aggressor privately observes the aggressor's offensive capability, θ_A, drawn according to the cumulative distribution function F^A, and chooses a demand a. Simultaneously, the defender privately observes the defender's defensive capability, θ_D, drawn from F^D, and offers d. Large values of θ_D and θ_A mean that the defender's defense is weak and the aggressor's offense is strong, which in turn means that the aggressor will do well if the conflict ends in war.

At the final stage 2, the payoffs are determined as follows:

(a) If $d \geq a$, then the conflict is settled at the average of the demand and offer, $[a + d]/2$.
(b) If $d < a$, then the demand and offer are inconsistent and A wars on D. The outcome of the war is then $[\theta_A + \theta_D]/2$.[5]

Each player also incurs two types of cost: (1) a fixed cost, $c \geq 0$, and (2) a variable cost with parameter $k \leq 1/2$, depending on the outcome of the war. For the defender, the variable cost is $k[\theta_A + \theta_D]/2$. That is, the worse the war outcome is for the defender, the higher the defender's variable cost incurred during the war. For the aggressor, k works in the opposite direction: $k - k[\theta_A + \theta_D]/2$. Thus when $[\theta_A + \theta_D]/2 = 0$, the cost to the aggressor will be $c + k$ and the cost to the defender will be $c + 0$; and when $[\theta_A + \theta_D]/2 = 1$, the cost to the aggressor will be $c + 0$ and the cost to the defender will be $c + k$.

Intuitive Explanation of the Major Results

As will be seen in the next section, deriving each belligerent's optimal bargaining strategy is quite complicated. It requires each side to calculate its best response to the other side's strategy, taking into account the cost incurred if there is a war and a recognition that each side has partial but different information regarding the outcome if a war takes place. Nevertheless, considerable intuition can be achieved without recourse to all of the math.

The aggressor's demand, a, is an upward sloping function of how well it expects to do in the war. Other things being equal, the greater its offensive strength (i.e., the larger θ_A is), the more it will demand. This can be seen in Figure 1, where the aggressor's demand curve, $a = A(\theta_A)$, is an upward sloping function of θ_A. Similarly, the defender's offer, $d = D(\theta_D)$, is an upward sloping function of how poorly its defensive strength is (the larger θ_D is, the weaker its defensive strength).

[5] I have implicitly assumed risk-neutral parties; so $[\theta_A + \theta_D]/2$ could be the *expected* outcome rather than *the* outcome. This value should be understood as the equilibrium outcome so that there will be no further impetus toward war until there is further technological change.

FIGURE 1 Demand and Offer Curves

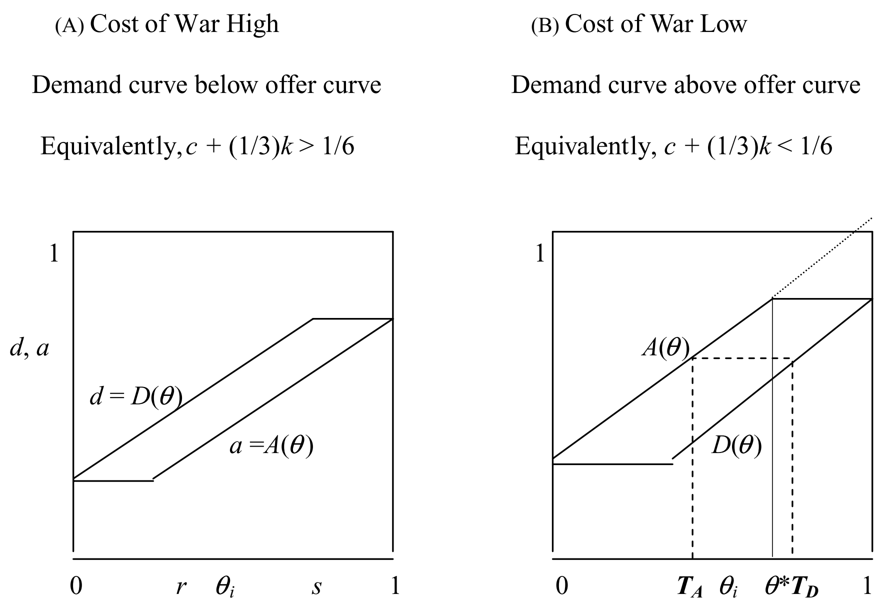

The minimum value of $d = D(\theta_d)$ is $2c - (1/6)[1 + k] + (1/2)k$ when $\theta_D = 0$ and the maximum value of d is $2c + (1/2)[1 + k] + (1/2)k$ when $\theta_D = 1$. So $a = A(\theta_A)$ is never strictly greater than the largest value of d nor strictly below the smallest value. Similarly, d is bounded by the natural range of a. The dashed line beyond θ^* is an inessential variation as the outcome will be the same.

Other things being equal, the higher the cost of war, the higher the offer curve by the defender and the lower the demand curve by the aggressor; that is, as c increases, the offer curve shifts upward and the demand curve shifts downward. Settlements are a substitute for war and the higher the cost of war, the more generous each side will be in bargaining. What is perhaps not so obvious is that if the cost of war is high enough, the desire to avoid war will result in the demand curve by the aggressor being below the offer curve of the defender as in Figure 1A.

A war occurs if and only if the demand is strictly greater than the offer. For example, in Figure 1A, if the defender observes r and the aggressor observes s, then there will be a war as $a = A(s) > D(r) = d$; while if the aggressor observes r and the defender observes s, then there will be a settlement as $A(r) < D(s)$. Ignore the kinks and the other items in Figure 1 until we get to the mathematical exposition.

We next turn our attention toward the likelihood of war. Recall that the outcome of the war is the average of the offensive capability by the aggressor and the defensive capability of the defender: $.5(\theta_A + \theta_D)$. Looking at Figure 1A, if the probability of the aggressor winning is very high relative to the median value, then *both* θ_A and θ_D must be very high.[6] Hence, the demand by the aggressor is less than the offer by the defender and a settlement takes place. Similarly, if the probability of the aggressor winning is very low, then *both* θ_A and θ_D must be very low. Again the demand is less than the offer and a settlement takes place. Next, consider the case where $.5(\theta_A + \theta_D) = .5$. One possibility is that $\theta_A < \theta_D$, in which case, there is again a settlement. And even if θ_A is slightly greater than θ_D, there will be a settlement. But if θ_A is much greater than θ_D, then a settlement will not take place. This is most easily illustrated for the extreme case where $\theta_A = 1$ and $\theta_D = 0$. *This suggests that when the cost of war is high and the demand curve is below the offer curve, the only time that we will have wars is in the middling-type cases* (*where* $5(\theta_A + \theta_D)$ is close to .5) *rather than at the extremes. In contrast, when the cost of war is low as in Figure 1B and the relationship between the demand and offer curves is reversed, the results are reversed.* Here, when there is a high probability of the aggressor winning (that is, $.5\theta_A + .5\theta_D$ is close to 1), it must mean that both θ_A and θ_D are close to 1. As a result the demand is greater than the offer and there is a war. The same holds if the probability of the

[6] Here, the language is in terms of probability, but it could also be in terms of the value of land captured by the aggressor.

aggressor winning is close to 0. Only toward the middle and only when θ_A is close to 0 and θ_D is close to 1 will there be a settlement. So in this case, *the most extreme situations are the least likely to settle.*

The Formal Derivation

In this section, I assume uniform signal distributions (as Chatterjee and Samuelson 1983; Filson and Werner 2002; Powell 1999 and others have done before). That is, $F^D(\theta) = F^A(\theta) = \theta$ and $dF^D(\theta) = dF^A(\theta) = d\theta$ for all θ in [0, 1]. The uniform distribution is very convenient because the objective functions take a simple form and the equilibrium conditions can be solved analytically.

The objective of the aggressor is to *maximize its expected return* (net of any war costs that might be incurred if the threat of war is carried out), conditioned on the aggressor's realized signal θ_A and the defender's strategy D. The payoff function for the aggressor is

$$\Pi^A(a, \theta_A, D, F^D) = 0.5 \int_{D^{-1}(a)}^{1} [a + D(x)]\, dx + 0.5 \int_{0}^{D^{-1}(a)} [(\theta_A + x)(1 + k) - 2c - 2k]\, dx. \quad (1)$$

Hence the first term represents the expected return to the aggressor (the value of the land that D cedes to A) due to a settlement and the second term represents the expected return to A when a settlement does not take place and A gains the land through war. Of course, when A goes to war, A incurs the cost of war, which is subtracted from the value of the land gained from war. As noted earlier, the better A does in the war, the lower A's variable cost.

A closer reading of equation (1) may be useful. The first term looks at those circumstances where there is a settlement; that is, where $d \geq a$. Now d depends on the defender's observation, θ_D, and the defender's bargaining strategy. This is the function $d = D(x)$. Although we use the standard notation, x, in fact x stands for some value of θ_D. Hence, we are looking at those values of θ_D such that $d = D(\theta_D) \geq a$; we can restate this in terms of the inverses: we are looking at those cases where $\theta_D \geq D^{-1}(a)$. So if there is to be a settlement, the term below the integral sign is the lower bound on θ_D.

Note that a settlement will be the average of the *demand and offer*; hence, we have $.5(a + D(x))$. Note also that we need to consider all $d \geq a$. Thinking in discrete terms, we find the probability that θ_D takes a certain value (the continuous cognate of this probability is dF^D) and multiply this times the settlement. We then sum this over all possible values of θ_D such that $d = D(\theta_D) \geq a$.

The second term considers those circumstances that lead to war (where $d = D(\theta_D) < a$). This time, it is the average of the *observations* that determines the outcome of the war; hence $.5(\theta_A + x)$, where once again x represents θ_D. Note that the d in dx is the derivative and has nothing to do with the offer by the defender.

In a nutshell, the aggressor's strategy must take into account the defender's strategy as well as the fact that both the aggressor and the defender are basing their strategies on a partial picture of the expected outcome of the war.

The aggressor's strategy A is a *best response* to the defender's strategy D. I write $A \in ABR(D)$ if, for each possible signal realization θ_A, the value $a = A(\theta_A)$ solves the problem $Max_y\, \Pi^A(y, \theta_A, D, F^D)$.

Similarly, the objective of the defender is to *minimize losses* (including any war costs that might be incurred), conditioned on the defender's realized signal θ_D and the aggressor's strategy A. The loss function for the defender is

$$\Pi^D(d, \theta_D, A, F^A) = 0.5 \int_{0}^{A^{-1}(d)} [d + A(y)]\, dy + 0.5 \int_{A^{-1}(d)}^{1} [(\theta_D + y)(1 + k) + 2c]\, dy. \quad (2)$$

Definition. A Nash equilibrium (*NE*) of the above game is a strategy pair (D, A) such that $A \in ABR(D)$ and $D \in DBR(A)$, the defender's best response to strategy A. More simply put, a Nash equilibrium exists when A's strategy response is best for A against D's strategy and D's strategy response is best for D against A's strategy.

I first show that the bargaining model has a pure-strategy equilibrium.

Proposition 1. The bargaining model has a Nash equilibrium in the piecewise-linear, continuous bid functions graphed in Figure 1. The functions are

$$a = A(\theta_A) = (2/3)\theta_A[1 + k] - 2c + (1/2)[1 + k] - (3/2)k, \quad (3)$$

truncated above at $min\{1, 2c + (1/2)[1 + k] + (1/2)k\}$ and below at $max\{0, 2c - (1/6)[1 + k] + (1/2)k\}$, and

$$d = D(\theta_D) = (2/3)\theta_D[1 + k] + 2c - (1/6)[1 + k] + (1/2)k, \quad (4)$$

truncated above at $min\{1, (7/6)[1 + k] - 2c - (3/2)k\}$ and below at $max\{0, -2c + (1/2)[1 + k] - (3/2)k\}$.

Formal proofs for Propositions 1–3 are in the appendix.

Not surprisingly, the better the aggressor expects to do if there is a war, the more the aggressor demands; similarly the better the defender expects the aggressor to do, the more the defender will offer. It can also be seen that the higher the fixed cost of war (c) is, the lower the demand by the aggressor and the higher the offer by the defender. The intuitive reason for θ_A and θ_D being preceded by a fraction is that each belligerent knows that its signal is only part of the truth and thus only partially responds to the signal.

Surprisingly, the aggressor's demand curve need not be above the defender's offer curve. That is, if both sides observe the same signal, the aggressor's demand need not be above the defender's offer. The aggressor's demand curve is below the defender's offer curve if and only if:

$$\begin{aligned} a &= (2/3)\theta[1+k] - 2c + (1/2)[1+k] - (3/2)k \\ &< (2/3)\theta[1+k] + 2c - (1/6)[1+k] + (1/2)k = d. \end{aligned} \tag{5}$$

This inequality holds when $(2/3)[1 + k] < 4c + 2k$ or $1/6 < c + (1/3)k$.

Figure 1B shows the more intuitive situation where, at *equal* signals, the aggressor demands more than the defender offers. The intuitive explanation for the aggressor demand curve being above the defender offer curve when the cost of war is low (and below the defender offer curve when the cost of war is high) is that the aggressor wants to receive as much as possible and the defender wants to pay as little as possible. Thus each side wants to extract the surplus for itself. When the cost of war is low, the issue of surplus extraction is paramount and thus the aggressor's demand curve is above the defender's offer curve. As the cost of going to war increases, each side will be more willing to settle. Although the one side's greater willingness to settle increases the intransigence of the other side, this effect is less than the direct effect of the increased cost on the other side. So the net effect of an increase in the cost of war is a shift downward of the aggressor's demand curve and a shift upward of the defender's offer curve.

We next turn our attention toward the truncations. The reason for the lower truncation in Figure 1A is that the aggressor has no incentive to demand less than the lowest possible offer from the defender. Demanding less would not increase the likelihood of a settlement (it is already 100%), but would reduce what the aggressor gains in a settlement. Likewise, the upper truncation reflects the fact that the defender has no incentive to offer more than the aggressor's highest demand.

One might ask whether there are other Nash equilibria. There are *trivial* Nash equilibria, where all conflicts end in war because both the aggressor and the defender make offers certain to be rejected, e.g., $A(\theta) = 1$ and $D(\theta) = 0$ for all signals θ.[7] There are also variations on strategies (3) and (4) that are *inessential* in that they induce the same outcomes (i.e., the same mapping from signals to payoffs). To illustrate, suppose that the aggressor's demand curve in Figure 1B were straight (so that it included the dashed line) instead of the kinked line. If the aggressor's signal were greater than θ^*, then both sides would go to war whether the aggressor's demand curve were straight or kinked (as the probability of the defender observing 1 is 0). The war outcome depends only on the signal, not the demand. Hence the outcome will be the same under both demand functions. Thus some of the truncations in Figure 1 are inessential. However, they keep the graphs of the functions within the unit square, which simplifies later calculations.

The uniqueness result covers equilibria that are symmetric in the sense that corresponds to the symmetry of information structure. For example, suppose that when the defender observes the signal 1/4, the defender offers 1/3. Symmetry would imply that when the aggressor observes 3/4, the aggressor demands 2/3. With this in mind, I make the following definition:

Definition. The strategies A and D of the bargaining model are symmetric if for all θ in [0,1] we have $A(\theta) = 1 - D(1-\theta)$ or, equivalently, $D(\theta) = 1 - A(1-\theta)$.

It is easy to see that the trivial Nash equilibrium strategies mentioned above are symmetric. So are the piecewise-linear strategies (3–4). Our next result is that the equilibrium is essentially unique in its class.

Proposition 2. All nontrivial piecewise-linear symmetric Nash equilibria of the bargaining model induce the same outcome as strategies (3–4).

In the rest of this section, "equilibrium" refers to the outcome generated by strategies (3–4). I will focus on the probability of war and the distribution of conflicts that go to war. The first comparative statics result shows that the probability of a war decreases as war cost c increases.

Proposition 3. The probability of war (weakly) decreases as c and/or k increases.

A quick look at equations (3) and (4) shows that the aggressor's demand (weakly) decreases and the defender's offer (weakly) increases as c increases. Intuitively, as the cost of war increases, the desire for each side to avoid war increases. Hence a negotiated settlement becomes more likely. I include the word "weakly," because for $k = 0$, the

[7] I will drop the subscripts from θ_A and θ_D when the meaning is clear.

probability of a war is 0 for $c \geq 1/3$. Parallel results to Proposition 3 occur in the literature. See, for example, Powell (1999) and Reed (2003).

Next, I compare the characteristics of conflicts that settle to those that end in war. Let $W = [\theta_A + \theta_D]/2$ be the expected outcome of the war given the private signals.

Proposition 4. The equilibrium probability of a war increases in $|W - 0.5|$ when $c + (1/3)k < 1/6$ and decreases in $|W - 0.5|$ when $c + (1/3)k > 1/6$. That is, the farther away the war outcome will be from the median war outcome, the more likely that there will be a war when war costs are low and the less likely there will be a war when war costs are high.

I have already provided the intuitive explanation. This proposition suggests that conflicts that are settled peacefully differ from conflicts that go to war in *certain* ways specified by the model. Because wars are not a random sample from the set of conflicts, one needs to be careful about making inferences from knowledge about war outcomes to expected outcomes for conflicts that are settled. Furthermore, wars with large c yield opposite conclusions from wars with small c.

Can wars occur when the belligerents are jointly pessimistic? If we interpret joint pessimism as θ_A being less than θ_D, then it is possible to have a war even if both sides are jointly pessimistic. This is captured in the following proposition.

Proposition 5. If c is low, then a war is possible even if the belligerents are jointly pessimistic (that is, $\theta_A < \theta_D$).

This can be immediately seen by referring back to Figure 1B. Choose any value for θ_A, which we will label $\boldsymbol{T_A}$. Then find the corresponding value of a on the aggressor's demand curve. Next, draw a horizontal line through a, and find where it intersects the defender's offer curve to the right. Finally, drop a vertical to discover the corresponding value of $\boldsymbol{\theta}_D = \boldsymbol{T_D}$. All values of θ_D strictly in between the bold terms, $\boldsymbol{T_A}$ and $\boldsymbol{T_D}$, will result in war even though the defender believes that the aggressor will do better than the aggressor believes to be the case. In a nutshell, because the demand curve is above the offer curve, a point modestly to the right on the offer curve will yield an offer less than the demand.

Comparison to Previous Theoretical Results

It is useful to compare the demand and offer curves derived here to those generated by other approaches.

As already noted there are many possible versions of the one-sided asymmetric information model of conflict involving different unknowns (outcome of war, cost to defender, resolve, etc.) with differing placement of the asymmetry (the defender may be the uninformed party or the aggressor may be the uninformed party), but all have a related mathematical structure. Only a few of these articles have considered which types of conflicts are most likely to result in war. So, I will derive the results here. I first consider the case where the defender draws θ_D, which is uniformly distributed on $[0, 1]$ and the outcome if there is a war is θ_D. The cost to each side is public information and, just to make the presentation simpler, cost is assumed to be the same for both sides. The aggressor makes a take-it-or-leave-it demand in the absence of any further information about the outcome of the war beyond the fact that the defender is drawing from the given distribution. The aggressor does not have a signal in this one-sided asymmetric model, but one could draw a horizontal line to represent the demand by the aggressor. The defender's offer curve is $c + \theta_D$; that is, the defender will accept any demand by the aggressor that is less than the total loss to the defender from going to war. Only when the defender has a low θ_D will the defender reject the aggressor's demand and the conflict ends in war. So this model predicts that wars sample from conflicts where the aggressor's gains are relatively small (but not when the aggressor's gains are relatively large).

Powell (1999) models the likelihood of war, which is the probability that one side will be dissatisfied times the probability that negotiations will be unsuccessful given that the aggressor is dissatisfied. In contrast, I only look at cases where one country is dissatisfied; so Powell is asking a somewhat different question. Nevertheless, I can make some comparisons as he also assumes a uniform distribution. Recall that in his model, the sides are unsure about each other's cost but not about the expected outcome of the war, $\bar{\theta}$. Hence, as long as the values of $\bar{\theta}$ and c are in the range so that all threats are credible (that is, the aggressor state will remain dissatisfied), which is the assumption made here, wars are a perfectly random selection from $\bar{\theta}$. On the other hand, if we allow for the possibility that both sides might be satisfied (as Powell does), then the most extreme cases are most likely to lead to war as the most extreme outcomes are the most likely for one side to be the most dissatisfied. In contrast, the model presented in this article predicts that when the cost of war is high, wars are concentrated at the center of potential expected war outcomes.

Because these different theoretical models predict different empirical relationships, the models can be subject to testing.

Extensions

Additional insights can be obtained by relaxing the assumptions of the basic model. First, consider shifting the distribution of war outcomes by adding a constant M so that the range of possible outcomes is $[0 + M, 1 + M]$. In the model, the probability of war depends on the differences in expectations, not on the level of these expectations. Hence such shifts in the outcome range have no effect, assuming that the threat of war is always credible (I will come to the issue of credibility shortly).

Next, consider increasing the width $(U - L)$ of the war outcome range to $(U + N) - (L - N)$, holding constant the cost $c \geq 0$, and maintaining the assumption that signals are independent and uniformly distributed on the interval. We have the following comparative statics result.

Proposition 6. An increase in the outcome spread increases the equilibrium probability of a war.

The formal proof is in the appendix. Intuitively, when the outcome spread increases, the relative cost of a war falls. And when the relative cost of a war falls, the probability of a war increases. Reed (2003) has a similar result for the one-sided information case, and I suspect that most existing models would produce an analogous result.

Relaxing the assumption of independently distributed signals does not appear to introduce any interesting new issues. The correlated signals can be decomposed into a common component that shifts the outcome midrange and private idiosyncratic components that are independent. As long as the belligerents are risk neutral, it seems that only the idiosyncratic components matter.

The ordinal properties of the main propositions are likely to hold even when the distribution is not uniform on the interval [L, H]. Recall the intuitive explanation for the results demonstrated in Figure 1. The intuition holds whenever the aggressor's demand curve is a shift upward from the defender's offer curve, not just when θ is distributed uniformly and the demand and offer curves are straight lines. I have assumed that both sides are risk neutral. It is therefore plausible to assume that on the margin both sides respond to an increase in the expected outcome of the war in the same way (that is, the curves have the same slope).

There is no reason to suppose informational asymmetries, but there is a natural strategic asymmetry: upon finding that the defender has offered less than the aggressor demanded, the aggressor might decide not to go to war. The model presented in the second section assumed that the aggressor would always go to war. I will now present the conditions where this assumption holds. To speed things along I will suppose that $k = 0$. The worst signal that the aggressor can observe is $\theta_A = 0$. Then according to equation (3), $a = 1/2 - 2c$. If the aggressor's demand is rejected, the aggressor can infer $(2/3)\theta_D + 2c - 1/6 < 1/2 - 2c$; i.e., $\theta_D < -6c + 1$. If the conflict goes to war, the aggressor expects to get $.5[0] + .5(.5)[1 - 6c] - c = .25 - 2.5c$. Hence $c < .1$ implies that the threat of a war is always credible in the basic model. For $c > .1$ in more general models, the threat of a war is still credible if $c < L$. Assuming $c < L$ is the standard approach used in the literature to ensure that the aggressor will always go to war if its demand is rejected.

Our setup is very versatile and can serve as the basis for more complicated endeavors. For example, Fearon's (1994) model of signaling and screening can be front-loaded onto the bargaining model presented here. To give a feel for how this could be done, suppose, as Fearon does, that there is an audience cost for the defending country. The willingness of the defending country to undertake an audience cost would then reveal to the aggressor that the defending country's observation, θ_D, is less than some number, say .8. Then the bargaining protocol as outlined in this article begins. The only difference is that the aggressor now knows that the defender's observation, θ_D, has been drawn from a uniform distribution on [0, .8] instead of on [0.1]. As a consequence, the demand by the aggressor is less than it would have been without the costly signal (if the signal is working as it is supposed to).

Empirical Tests

War and misperception are inextricably bound. . . .
Morrow (1989, 954)

This article has provided a number of empirical hypotheses. An obvious set of questions to ask is whether there is any extant evidence that supports (or contradicts) the theory, and if not, how one might go about getting the requisite data and testing the hypotheses.

I start with Proposition 6, which states that, other things being equal, the higher the variance in $.5(\theta_A + \theta_D)$, the greater the likelihood of war. In general, it is hard to find a direct measure of variance (the degree of uncertainty that each side has about the other side's private information). So researchers try to find correlates to variance, although the rationale for such correlation may at times be tenuous. One way of interpreting the power parity explanation for war is that when there is putative parity in power, it is unclear where the balance lies; hence there is greater uncertainty about relative power compared to a situation where there is a preponderance of power and therefore power parity produces a greater likelihood of war (see Organski and Kugler 1980, and Kugler and Lemke 1996). Reed (2003) provides both a model of misperception and statistical support. However, despite his valiant efforts, other things are not equal when one compares power parity to preponderance of power; in particular, the cost of war, which we have shown to be critical for our understanding, may vary both within and across these two types.[8] One can also look for factors that reduce variance. For example, improvements in spying should reduce uncertainty about the other side's private information, thereby reducing the probability of war.

In Proposition 3, I noted that an increase in the cost of war reduces the likelihood of war. This may explain why we did not see any direct wars between the superpowers, which in all likelihood would have involved nuclear weapons and great cost to both sides although there were many relatively low-cost proxy wars between them.

If we want to get to more subtle empirical issues, we encounter a problem that exists for all bargaining models, not just the one presented here: regardless of whether the private information is one-sided or two-sided or whether the lack of information is about the opposing side's strength, cost of war, or resolve—war is in the error term.[9] That is, it is the differential between each side's *private* information and the other side's expectation that drives the parties to war (for example, θ_D is much lower than the aggressor expects it to be and/or θ_A is much higher than the defender expects it to be). Hence, we have the quote at the beginning of the section about war and misperception being inextricably bound. Notice that the word "private" is italicized. This is to emphasize the point that private information is not publicly available at the time of war; otherwise, there would have been no disagreement about the outcome between the two parties and a settlement would have occurred (assuming that agreements are upheld). Nor is such information typically available to researchers ex post. At a minimum, the appropriate methodology requires some measure of expectations, a variable that is not usually included in empirical studies of war, and, in any event, is not as easily obtained as other data on bargaining and war.

I will consider two potential ways of getting a handle on the expectations variable. The first method works hand in hand with the methodology of this article. Before the bargaining starts, each country knows that the other side is drawing from a distribution bounded by [L, H], with expected value $(L + H)/2$.[10] Present-day research studies based on publicly known data at the time (e.g., relative size of the armies) could estimate L, H and $(L + H)/2$. It is reasonable to assume that these are unbiased estimates.

To see how this fits in with our previous analysis, assume that the researcher has the same information about the defender that the aggressor has about the defender, and similarly that the researcher has the same information about the aggressor as the defender has about the aggressor. Then the researcher's predicted outcome if a war should happen would be $.5(E[\theta_A] + E[\theta_D]) = .5[(L + H)/2 + (L + H)/2] = .5(L + H)$. Suppose first that the cost of war were *very* high so that the offer curve is above the demand curve as in Figure 1A.[11] If θ_A and θ_D were both high, both low, or both near the expected value, then there would be an agreement and no war would take place. Only if θ_A were high and θ_D were low would there be war and in such an event, the outcome of the war would be relatively close to the researcher's predicted outcome, $.5(L + H)$, although the constituent parts would not be close to their expected values. Suppose next that the cost of war were very low. If θ_A and θ_D were both high, both low, or both near the expected value, then there would not be an agreement and a war would take place (see Figure 1B).[12] Only if θ_A were relatively low and θ_D were relatively high would there not be a war. So, when the cost of war is low, the outcome of the war is likely to be greatly different from the researcher's predicted outcome (one way or another) in comparison to the case where the cost of war is very high.[13] In a nutshell, *one should find larger discrepancies one way or the other between predictions based on public knowledge and the outcome of the war when the cost of the war is low than when the cost of the war is high.* In general, each side is likely to have more information about the other side than the researcher does. So large errors and the probability of war are less likely than the researcher might predict, but the logic as outlined above still holds.

[8]See also Wagner (1994) and Powell (1999), who show that the threat of war is less likely to be credible when there is parity in power than when there is preponderance of power.

[9]See Gartzke (1999) for an extensive discussion.

[10]If $[H, L] = [0, 1]$, then the expected value, E, of the other side's draw is .5.

[11]Note that the word "high" has a precise meaning (see Proposition 4).

[12]Note that war reveals the average of the private signals.

[13]The θ_i can be viewed as unbiased estimates so that $.5(\theta_A + \theta_D)$ is itself an unbiased estimate of the outcome of the war. Hence, the relationships are stochastic rather than certain. As already noted, high-cost wars are less likely to occur.

Now the nice part about this methodology is that we need not only look at wars. We can look at settlements, also. If we focus on settlements, all the relationships are reversed. If wars are high cost, then settlements may vary greatly in one direction or the other from the researcher's predicted estimate of the settlement, which in turn is the researcher's predicted outcome of the war if it would have taken place. While if wars are low cost, settlements will tend to be much closer to the researcher's estimate of the settlement.

In the absence of a systematic statistical study devoted to the question at hand, one can only find illustrations rather than conclusive evidence. With this disclaimer in mind, I will provide an illustration. Austria was expected to win the Seven Week War with Prussia. As Bueno de Mesquita (1990) has shown, this was reflected in historical accounts at the time, financial markets, and present-day research making use of *The Correlates of War* data set. The outcome of this low-cost war was quite different from expectations (again reflected in financial markets). Prussia prevailed and that was the beginning of the end of the Austrian empire.

The second method requires extensive archival work. To explain why a particular war did or did not happen, one must discover the views of each belligerent regarding their own and their adversary's strengths and weaknesses (including costs) and parameterize these expectations. In this way one discovers the values of θ_A and θ_D as well as each side's estimates of the other side's parameters.

Sometimes new theoretical results lead to new empirical developments and sometimes the direction of causation is in the opposite direction. This article is in line with the former possibility. The material in this section provides a more direct empirical strategy that should be undertaken for all bargaining models and at the same time suggests a set of questions to be empirically tested in the future given the particular model presented here.

Concluding Remarks

I have modeled bargaining in the shadow of war when there is two-sided incomplete information with common values. In particular, both the percentage of land captured (more generally, the extent of policy concessions made by the defender) and the cost of the war to each side depend on both belligerents' private signals. I demonstrated the existence of an equilibrium and derived a closed-form solution for the uniform distribution case. I also derived a number of comparative-static results. Some of these results are intuitive: (1) higher war costs reduce the likelihood of war; (2) greater variance in the set of possible outcomes increases the likelihood of war (even though the participants are risk neutral); and (3) when war costs are high (and war is a credible threat), those conflicts with middling war outcomes are less likely to be settled than those conflicts where the war outcome is likely to be at one extreme or the other. Other results are counterintuitive: (1) the aggressor's offer curve may be below the defender's demand curve; (2) war may occur even if both sides are jointly pessimistic about the outcome; and (3) when war costs are low, those conflicts with middling war outcomes are more likely to be settled than those conflicts where the outcome is likely to be at one extreme. These latter results differ from those of earlier models.

The model suggests why earlier empirical studies on whether balance of power or preponderance of power is more likely to lead to war have been contradictory. These empirical studies have not controlled for all of the relevant variables (in particular, the relative cost of war). So the hard job of doing empirical research will be even harder in the future.

I have shown how the basic model can be extended to more general situations where the distribution need not be uniform and the bargaining need not be symmetric. Perhaps most important, I was able to encapsulate the logic of the situation into a simple diagram so that the logic can be grasped by those without sophisticated mathematical skills. Supply and demand curves are an important tool for understanding economic markets, even for those who have only a vague understanding of utility maximization and profit functions. In a similar way, the demand and offer curves presented here can be used for understanding conflict, even for those who have only a vague understanding of the derivation of these curves. Just as supply and demand curves are widely used in economics, I hope that the demand and offer curves presented in this article can be similarly and as widely used in international relations and conflict studies.

Appendix

Proposition 1. The bargaining model has a Nash equilibrium in the piecewise-linear, continuous bid functions graphed in Figure 1. The functions are

$$A(\theta_A) = (2/3)\theta_A[1 + k] - 2c + (1/2)[1 + k] - (3/2)k \tag{A1}$$

truncated above at $min\{1, 2c + (1/2)[1 + k] + (1/2)k\}$ and below at $max\{0, 2c - (1/6)[1 + k] + (1/2)k\}$, and

$$D(\theta_D) = (2/3)\theta_D[1 + k] + 2c - (1/6)[1 + k] + (1/2)k, \tag{A2}$$

truncated above at $min\{1, (7/6)[1 + k] - 2c - (3/2)k\}$ and below at $max\{0, -2c + (1/2)[1 + k] - (3/2)k\}$.

Proof. We need only verify that (A1) is a best response to (A2) and vice versa. Note first that it is impossible for both the demand and offer to equal 1. $A(\theta_A) = (2/3)\theta_A[1 + k] - 2c + (1/2)[1 + k] - (3/2)k \geq 1$ only if $c \leq 1/12 - (1/6)k$.[14] But if this is the case, then $D(\theta_D) = (2/3)\theta_D[1 + k] + 2c - (1/6)[1 + k] + (1/2)k \leq (2/3)[1+ k] + 1/6 - (1/3)k - (1/6)[1 + k] + (1/2)k = (2/3)[1 + k]$. The last term is less than 1 for $k < \frac{1}{2}$. A similar exercise shows that both the demand and offer cannot equal 0. Thus the 0, 1 truncations are inessential because when the truncations are operative, the conflict would go to war whether there was a truncation or not, and the demands and offers do not affect the war outcome. The truncation of $A(\theta_A)$ above at $2c + 1/2 + k$ and the truncation of $D(\theta_D)$ below at $-2c + 1/2 - k$ are also inessential for similar reasons.

The defender minimizes its expected payment

$$\Pi^D(d, \theta^D, A, F^A) = 0.5\int_0^{A^{-1}(d)} [d + A(y)]\,dy + 0.5\int_{A^{-1}(d)}^1 [(\theta_D + y)(1 + k) + 2c]\,dy. \tag{A3}$$

Note that $A^{-1}(d)$ is that θ_A such that $A(\theta_A) = d$ and that $A(A^{-1}(d)) = d$.

Differentiating this expression (A3) with respect to d, one obtains the first-order condition:

$$\Pi^D_d = d/A'(A^{-1}(d)) + .5A^{-1}(d) - [.5A^{-1}(d)(1 + k) + .5\theta_D(1 + k) + c]/A'(A^{-1}(d)) = 0. \tag{A4}$$

Multiplying (A4) through by $A'(A^{-1}(d))$, we get the more convenient expression:

$$d + .5A^{-1}(d)A'(A^{-1}(d)) - .5\theta_D(1 + k) - c - .5A^{-1}(d)(1 + k) = 0. \tag{A5}$$

We first look at the conflict where the aggressor's demand curve is above the defender's offer curve and the truncations are inessential; so we can ignore them. $A(\theta)$ is above $D(\theta)$ when $A(\theta) = (2/3)\theta[1 + k] - 2c + (1/2)[1 + k] - (3/2)k > (2/3)\theta[1 + k] + 2c - (1/6)[1 + k] + (1/2)k = D(\theta)$. Equivalently, when $-2c + (1/2)[1 + k] - (3/2)k > 2c - (1/6)[1 + k] + (1/2)k$, or $2/3 > 4c + (4/3)k$. Equivalently, $A(\theta)$ is above $D(\theta)$ when $1/6 > c + (1/3)k$.

We are checking the defender's best response to $A(\theta_A) = (2/3)\theta_A(1 + k) - 2c + (1/2)[1 + k] - (3/2)k$, so we substitute[15] $A^{-1}(d) = (3/2)d/[1 + k] + 3c/[1 + k] - (3/4) + (9/4)k/[1 + k]$ and $A' = (2/3)[1 + k]$ into (A5) to obtain

$$\begin{aligned}
&d + .5\big[(3/2)d/[1 + k] + 3c/[1 + k] - (3/4) + (9/4)k/[1 + k]\big](2/3)[1 + k] - .5\theta_D[1 + k] - c \\
&\quad - .5\big[(3/2)d/[1 + k] + 3c/[1 + k] - (3/4) + (9/4)k/[1 + k]\big][1 + k] \\
&= d + .5\big[d + 2c - (2/4)[1 + k] + (3/2)k\big] - .5\theta_D[1 + k] - c - .5\big[(3/2)d + 3c - (3/4)[1 + k] + (9/4)k\big] \\
&= d + d/2 + c - 1/4[1 + k] + (3/4)k - .5\theta_D[1 + k] - c - (3/4)d - (3/2)c + (3/8)[1 + k] - (9/8)k \\
&= (3/4)d - (3/2)c + (1/8)[1 + k] - .5\theta_D(1 + k) - (3/8)k = 0.
\end{aligned} \tag{A6}$$

Equivalently,

$$d = 2c - (1/6)[1 + k] + (2/3)\theta_D[1 + k] + (1/2)k. \tag{A7}$$

The unique solution is $d = (2/3)\theta_D\,[1 + k] + 2c - (1/6)[1 + k] + (1/2)k$, as desired. Next, we will find the second derivative. First, recall that $A' = (2/3)[1 + k]$; so $A''(A^{-1}(d)) = 0$. The second derivative of the objective function is therefore

$$\begin{aligned}
\Pi^D_{dd} &= 1.5/A'(A^{-1}(d)) - .5(1 + k)/[A'(A^{-1}(d))]^2 \\
&= 1.5/[(2/3)[1 + k]] - .5/[(4/9)(1 + k)] \\
&= [9/4 - (9/8)]/[1 + k] \\
&> (9/8)/[1 + k] > 0.
\end{aligned} \tag{A8}$$

So we are indeed at a minimum. Hence we have verified the best response for $c + (1/3)k < 1/6$.

The aggressor maximizes its net expected return

$$\Pi^A(a, \theta^A, D, F^D) = 0.5\int_{D^{-1}(a)}^1 [a + D(x)]\,dx + 0.5\int_0^{D^{-1}(a)} [(\theta_A + x)(1 + k) - 2c - 2k]\,dx \tag{A9}$$

Differentiating this expression with respect to a, one obtains the first order condition:

[14] When $\theta_A = 1$, then $a = (2/3)[1 + k] - 2c + (1/2)[1 + k] - (3/2)k = 7/6 - (2/6)k - 2c \geq 1$. Equivalently, $1/12 - (1/6)k \geq c$.

[15] We set $d = (2/3)\theta_A(1 + k) - 2c + (1/2)[1 + k] - (3/2)k$. Equivalently, $d + 2c - (1/2)[1 + k] + (3/2)k = (2/3)\theta_A(1 + k)$. Multiplying through by $(3/2)/[1 + k]$, we get $(3/2)d/[1 + k] + 3c/[1 + k] - (3/4) + (9/4)k/[1 + k] = \theta_A$.

$$\Pi_a^A = -a/D'(D^{-1}(a)) + .5 - .5D^{-1}(a) + [.5D^{-1}(a)(1+k) + .5\theta_A(1+k) - c - k]/D'(D^{-1}(a)) = 0. \quad \text{(A10)}$$

Multiplying (A10) through by $D'(D^{-1}(a))$, we get the more convenient expression:

$$-a + [.5 - .5D^{-1}(a)]\, D'(D^{-1}(a)) + [.5D^{-1}(a)(1+k) + .5\theta_A(1+k) - c - k] = 0. \quad \text{(A11)}$$

Recall that $D(\theta_D) = (2/3)\theta_D[1 + k] + 2c - (1/6)[1 + k] + (1/2)k$. So, $D' = (2/3)[1 + k]$, and

$$D^{-1} = (3/2)a/[1+k] - 3c/[1+k] + (1/4) - (3/4)k/[1+k]. \quad \text{(A12)}$$

Hence

$$\begin{aligned}
\text{(A11)} &= -a + \big[.5 - .5[(3/2)a/[1+k] - 3c/[1+k] \\
&\quad + 1/4 - (3/4)k/[1+k]\big](2/3)[1+k] \\
&\quad + .5\big[(3/2)a/[1+k] - 3c/[1+k] \\
&\quad + 1/4 - (3/4)k/[1+k]\big](2/3)[1+k] \\
&\quad + .5\theta_A(1+k) - c - k \\
&= -a + .5 - .5[a - 2c + (2/12)[1+k] - (2/4)k \\
&\quad + .5\big[(3/2)a - 3c + (1/4)[1+k] - (3/4)k\big] \\
&\quad + .5\theta_A(1+k) - c - k \\
&= -a + (1/3)[1+k] - (1/2)a + c \\
&\quad - (1/12)(1+k) + (1/4)k + (3/4)a - (3/2)c \\
&\quad + (1/8)[1+k] - (3/8)k + .5\theta_A[1+k] - c - k \quad \text{(A13)} \\
&= -(3/4)a - (3/2)c + (3/8)[1+k] \\
&\quad + .5\theta_A[1+k] - (9/8)k = 0. \quad \text{(A14)}
\end{aligned}$$

Equivalently,

$$a = (2/3)\theta_A[1+k] - 2c + (1/2)[1+k] - (3/2)k. \quad \text{(A15)}$$

This is (A1).

Next suppose that $c + k/3 \geq 1/6$, so that the defender offer curve is on or above the aggressor demand curve and the demand and offer functions involve essential truncations as in Figure 1A. As explained in the text, the truncation does not reduce the probability of a settlement, but does make the settlement more favorable for the belligerent. The computations above for $c + (1/3)k < 1/6$ still hold, but we still need to check the truncations. It suffices to show that even when $\theta_D = 0$, the defender will still want to settle.

When $\theta_D = 0$, the defender's offer is $d = 2c - (1/6)[1 + k] + (1/2)k$. For there to be a settlement, the aggressor's demand must be less than or equal to $2c - (1/6)[1 + k] + (1/2)k$. If $2c - 1/6[1 + k] + (1/2)k > 0$, then there will be a *truncation* in the aggressor's demand curve and a settlement will take place when $2c - 1/6[1 + k] + (1/2)k \geq (2/3)\theta_A[1 + k] - 2c + (1/2)[1 + k] - (3/2)k$. Equivalently, when $4c - 2/3[1 + k] + 2k = 4c - 2/3 + (4/3)k \geq (2/3)\theta_A[1 + k]$. Equivalently, when $[6c - 1 + 2k]/[1 + k] \geq \theta_A$, a settlement will take place.

We will now establish that there is a Nash equilibrium when the aggressor truncates the demand curve. We have already demonstrated the aggressor's best strategy. We next show that the defender will not want to change its strategy even though the aggressor's strategy now has a kink in it. If the defender reduced its minimal offer ever so slightly, then there would no longer be a settlement when $[6c - 1 + 2k]/[1 + k] \geq \theta_A$ and $\theta_D = 0$. This would make the defender worse off. The logic is as follows. Suppose that $a = [6c - 1 + 2k]/[1 + k]$. Then it is the case that the aggressor's signal is uniformly distributed between 0 and $[6c - 1 + 2k]/[1 + k]$. Hence, the expected cost of the war to the defender would be: $\big[.5(0) + .5(.5)[6c - 1 + 2k]/[1 + k]\big][1 + k] + c = (5/2)c - 1/4 + (1/2)k$.

Let us first look at the term in the large brackets. Recall that the expected cost of the war depends on the average of the two observations by A and D. That is why there is a .5 before 0 and a .5 before (.5). Given that θ_A is uniformly distributed on $\{0, [6c - 1 + 2k]/[1 + k]\}$, the expected value of θ_A is half that much (thus the third .5). All of this is multiplied times $[1 + k]$ in order to account for the variable cost of war. Finally, the fixed cost, c, is added.

If $\theta_D = 0$, when is the expected cost of a war greater than the cost of a settlement? When $(5/2)c - 1/4 + (1/2)k > 2c - (1/6)[1 + k] + (1/2)k$. Equivalently, when $(1/2)c + (1/6)k > 1/12$, or $c + (1/3)k \geq 1/6$. But this is our condition for the aggressor's demand curve to be above the defender's offer curve. So D will want to settle rather than lower its offer. If the defender were to raise its offer, the defender would clearly be worse off.

The game being symmetric, a similar argument verifies that the given piecewise-linear aggressor strategy is a best response to the defender's given strategy.

Proposition 2. All nontrivial piecewise-linear symmetric Nash equilibria of the bargaining model induce the same outcome as strategies (3–4).

Proof. We will now make use of the symmetry conditions to rewrite (A5) as a function of D rather than A. We first explicitly consider the following symmetry relationships.

Let $y = A(z) = 1 - D(1 - z)$. Then $z = A^{-1}(y)$ and $D^{-1}(1 - y) = 1 - z$. Equivalently, $z = 1 - D^{-1}(1 - y)$.

So $A^{-1}(y) = 1 - D^{-1}(1 - y)$ and $A'(A^{-1}(d)) = D'(1 - A^{-1}(d)) = D'(1 - 1 + D^{-1}(1 - d)) = D'(D^{-1}(1 - d))$.

Substituting these relationships into (A5) we get:

$$\begin{aligned} &d + .5[1 - D^{-1}(1 - d)]\, D'(D^{-1}(1 - d)) \\ &\quad - .5\theta_D - c - .5[1 - D^{-1}(1 - d)] \\ &= d + .5[1 - D^{-1}(1 - d)] \\ &\quad \times [D'(D^{-1}(1 - d)) - 1] - .5\theta_D - c = 0 \quad \text{(A16)} \end{aligned}$$

D is assumed to be piecewise-linear. Let $d = a_d + b_d c + e_d \theta_D$, then $\theta_D = (d - a_d - b_d c)/e_d$ and $D' = e_d$. Equation (A8) can be rewritten as follows:

$$\begin{aligned} &d + .5(e_d - 1) - .5(e_d - 1) \\ &\quad \times (1 - d - a_d - b_d c)/e_d - .5\theta_D - c = 0. \quad \text{(A17)} \end{aligned}$$

Equivalently,

$$\begin{aligned} &d(1.5e_d - .5)/e_d = -.5(e_d - 1)(1 - 1/e_d) \\ &\quad - .5(e_d - 1)a_d/e_d - .5(e_d - 1)(b_d c)/e_d + c + .5\theta_D, \quad \text{(A18)} \end{aligned}$$

or

$$\begin{aligned} d = &-.5(e_d - 1)(1 - 1/e_d)e_d/(1.5e_d - .5) \\ &- .5(e_d - 1)a_d/(1.5e_d - .5) - .5(e_d - 1)(b_d c)/ \\ &(1.5e_d - .5) + ce_d/(1.5e_d - .5) \\ &+ .5\theta_A e_d/(1.5e_d - .5). \quad \text{(A19)} \end{aligned}$$

Now the coefficient of θ_D is e_d. So from (A19), we have the following relationship:

$$e_d = .5e_d/(1.5e_d - .5) = e_d/(3e_d - 1). \quad \text{(A20)}$$

The solutions are $e_d = 2/3$, 0. Clearly, the pair (A1–A2) satisfies these conditions. The inessential truncations need not hold because, as shown in Proposition 1, the probability of a war is 1 regardless. So the outcome is not changed as long as the slope is greater than or equal to zero once the point of truncation is reached (hence the word "piecewise").

The only question remaining is whether there is another piecewise-linear function with this set of slopes but in a different combination. We first show that $A = 1/2 + 0\theta_A$, $D = 1/2 + 0\theta_D$ is not a Nash equilibrium for all $0 < c < 1/3$. Suppose that $\theta_A = 1$, then the expected outcome if the conflict goes to war is 3/4. So the aggressor will raise its demand above .5 if $c + .25k < .25$.

Every other horizontal line where D is below A always results in war and therefore is trivial. The reverse can never be an equilibrium as all conflicts would be settled and the aggressor would want to increase its demand and the defender would want to decrease its offer.

There are two other families of possibilities:

(1) The 2/3-slope line is broken up by one or more 0 slope lines.
(2) The 2/3-slope line is still in the middle, but the 0 slope line(s) start or stop at a different place.

We will focus on (1). The argument for (2) is a blend of the previous arguments.

Let us consider a horizontal portion between two line segments with 2/3 slope. Moving along the horizontal portion, the defender's loss from going to war strictly increases continuously as θ_D increases. Furthermore, we know that the probability of war is less than 1 as the defender increases its offer for still larger values of θ_D to reduce the probability of war. Thus, the defender should continuously strictly increase its offer to continuously reduce the probability of a war. But a 0 slope says otherwise. Hence we are led to a contradiction.

Proposition 3. The probability of war (weakly) decreases as c and/or k increases.

Proof. A conflict results in war if and only if

$$\begin{aligned} a &= (2/3)\theta_A[1 + k] - 2c + (1/2)[1 + k] - (3/2)k \\ &> (2/3)\theta_D[1 + k] + 2c - (1/6)[1 + k] + (1/2)k = d. \end{aligned}$$

That is, $(2/3)[1 + k]\,[\theta_A - \theta_D] > 4c - (2/3)[1 + k] + 2k$.

Equivalently, a conflict results in war if and only if

$$\begin{aligned} \theta_A - \theta_D &> 6c/[1 + k] - 1 + 3k/[1 + k] \\ &= [6c - 1 + 2k]/[1 + k]. \end{aligned}$$

Clearly, as c increases the probability of a war (weakly) decreases.

Next we consider the effect of an increase in k. Taking the derivative of the expression to the right of the equality with respect to k, we get $2/[1 + k] - [6c - 1 + 2k]/[1 + k]^2 = [-6c + 3]/[1 + k]^2$. For $c < 1/3$, an increase in k will also decrease the probability of war. For $c > 1/3$, the probability of war is 0 regardless of the value of k.

Proposition 6. An increase in the outcome spread increases the equilibrium probability of a war.

Proof. Consider increasing the width $U - L$ of the war outcome range to $(U + N) - (L - N)$, holding constant the cost $c \geq 0$ and maintaining the assumption that signals are independent and uniformly distributed. It is easy to check that independent uniform signal distributions retain those properties under the transformation. Note that the transformed war cost is $C = c\,(U - L)/(2N + U - L)$.

Since the normalized cost decreases as N increases from 0, we apply Proposition 3 to conclude that the equilibrium war probability increases.

References

Bueno de Mesquita, Bruce. 1990. "The Origins of German Hegemony." *World Politics* 43(October): 28–52.

Chatterjee, Kalyan. 1981. "Comparison of Arbitration Procedures: Models with Complete and Incomplete Information." *IEEE Transactions on Systems, Man and Cybernetics* 11(2): 101–9.

Chatterjee, Kalyan, and William Samuelson. 1983. "Bargaining under Incomplete Information." *Operations Research* 31(5): 835–51.

Fearon, James D. 1994. "Test of a Crisis Bargaining Model." *Journal of Conflict Resolution* 38(June): 236–69.

Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3): 379–414.

Filson, Darren, and Suzanne Werner. 2002. "A Bargaining Model of War and Peace: Anticipating the Onset, Duration, and Outcome of War." *American Journal of Political Science* 46(4): 819–38.

Gartzke, E. 1999. "War Is in the Error Term." *International Organization* 53(Summer): 567–87.

Kugler, Jacek, and Douglas Lemke. 1996. *Parity and War: Evaluations and Extensions of the War Ledger*. Ann Arbor: University of Michigan Press.

Morrow, James D. 1989. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33(4): 941–72.

Organski, A. F. K., and Jacek Kugler. 1980. *The War Ledger*. Chicago: University of Chicago Press.

Powell, Robert. 1999. *In the Shadow of Power*. Princeton, NJ: Princeton University Press.

Powell, Robert. 2002. "Bargaining Theory and International Conflict." *Annual Review of Political Science* 5: 1–30.

Powell, Robert. 2004. "Bargaining and Learning While Fighting." *American Journal of Political Science* 48(2): 344–61.

Reed, William. 2003. "Information, Power, and War." *American Political Science Review* 97(4): 628–41.

Reiter, Dan. 2003. "Exploring the Bargaining Model of War." *Perspectives on Politics* 1(March): 1–17.

Slantchev, Branislav L. 2003. "The Principle of Convergence in Wartime Negotiations." *American Political Science Review* 97(4): 621–32.

Wagner, Harrison. 1994. "Peace, War and the Balance of Power." *American Political Science Review* 88(3): 593–607.