Annu. Rev. Polit. Sci. 2002. 5:1–30
DOI: 10.1146/annurev.polisci.5.092601.141138

BARGAINING THEORY AND INTERNATIONAL CONFLICT

Robert Powell
Department of Political Science, 220 Barrows Hall, University of California at Berkeley, Berkeley, California 94720-1950; e-mail: rpowell@socrates.berkeley.edu

Key Words causes of war and conflict, inefficiency and breakdown

■ **Abstract** International relations theory has long seen the origins, conduct, and termination of war as a bargaining process. Recent formal work on these issues draws very heavily on Rubinstein's (1982) seminal analysis of the bargaining problem and the research that flowed from it. There is now what might be called a standard or canonical model of the origins of war that sees this outcome as a bargaining breakdown. This essay reviews this standard model and current efforts to extend it to the areas of (*a*) multilateral bargaining, which is at the heart of old issues such as balancing and bandwagoning as well as newer ones such as the role of third-party mediation; (*b*) the effects of domestic politics on international outcomes; (*c*) efforts to explicitly model intra-war bargaining; and (*d*) dynamic commitment problems.

INTRODUCTION

Bargaining—be it over the terms of a peace settlement, an alliance, a treaty, a trade agreement, or the structure of an international institution—is at the center of many of the most important issues in international politics. Not surprisingly, then, international relations theory has often looked to bargaining theory.[1] This is especially true of the most recent formal work on the origins, conduct, and termination of war, which draws very heavily on Rubinstein's (1982) seminal analysis and the research that flowed from it. Grounded in bargaining theory and building on earlier formal and nonformal analyses of war, the latest efforts are maturing into a coherent and cumulating body of research with well-defined questions; clear, deductive analyses; and empirically testable hypotheses. Most of this work is still largely theoretical, as might be expected in this relatively early stage in the development of this latest wave of research. Some testing has already been done,

[1]Conversely, the study of international politics seems to have stimulated important work in bargaining theory, most notably Schelling's *The Strategy of Conflict* (1960).

but the challenge for the future—as with so much of the broad thematic work in international relations—is to conduct compelling empirical tests while continuing to develop the theory.

This essay reviews the theoretical work on bargaining and war. The next section surveys results derived from noncooperative bargaining theory. Subsequent sections describe the basic bargaining-problem framework for studying war and its application to four areas: (*a*) multilateral bargaining, which is being used to study old issues such as balancing and bandwagoning as well as newer ones such as the role of third-party mediation; (*b*) the effects of domestic politics on international outcomes; (*c*) efforts to explicitly model intra-war bargaining; and (*d*) dynamic commitment problems.[2]

AN OVERVIEW OF NONCOOPERATIVE BARGAINING THEORY

Bargaining is about deciding how to divide the gains from joint action. That is, coordinated action frequently increases the size of the "pie"—for example, the exchange of goods often creates gains from trade; revising the territorial status quo peacefully rather than through the costly use of force means that the resources that would have been destroyed by fighting can now can be divided. The existence of potential gains from acting jointly creates an incentive to cooperate. But, of course, each actor also wants to maximize its share of those gains and, indeed, may take steps that reduce the chances of agreement when such steps promise a sufficiently large share of the gains if there is an agreement.

In 1982, Rubinstein's striking analysis renewed interest in studying bargaining with noncooperative game theory. The noncooperative approach focuses on the setting in which the negotiations take place and on how that setting shapes the bargaining strategies and ultimate outcomes. In particular, it models the bargaining problem as a noncooperative game and characterizes the equilibria of this game. Once this is done, changes in the bargaining setting are modeled by changing the underlying game and then tracing the effects of these changes on the game's equilibria. This section briefly reviews Rubinstein's (1982) analysis and some of the work that grew out of it [see Fudenberg & Tirole (1991, pp. 397–434), Kennan & Wilson (1993), and Muthoo (1999) for more extensive

[2]Space limitations preclude the discussion of important work on war (e.g., Fearon 1994, Kydd 1997, Downs & Rocke 1994) that is not based on bargaining models. Although the distinction is somewhat arbitrary, bargaining models give players a significant range of options when deciding how much of the bargaining surplus to demand. This contrasts with, for example, a war of attrition in which each player must demand everything or give in.

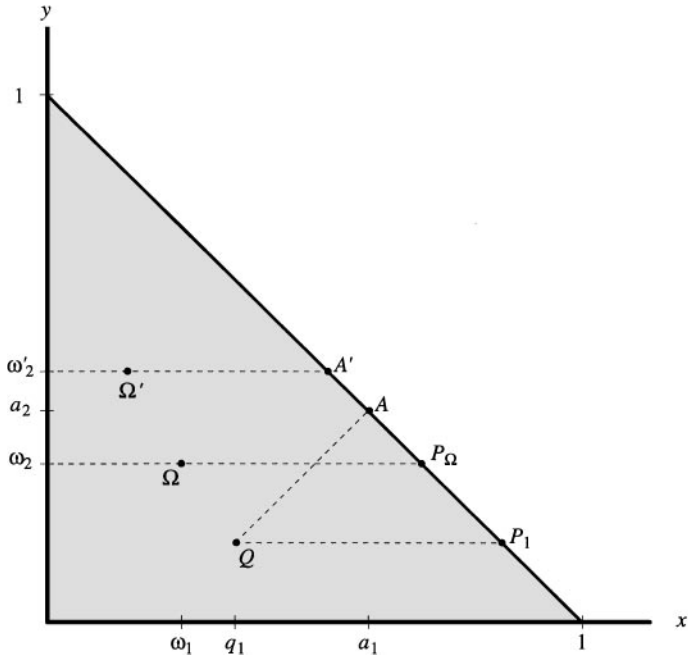

Figure 1 The bargaining problem.

reviews]. Special emphasis is given to the bargaining structures that have been used in applied work, and the outcomes these structures typically induce.[3] Suppose two players, *1* and *2*, are bargaining about how to divide the gains from cooperation. The shaded region in Figure 1 depicts the set of feasible outcomes and payoffs. For expositional simplicity, the bargainers are assumed to be risk-neutral, which means that the players' utilities to agreeing to (x, y) are, respectively, $U_1(x) = x$ and $U_2(y) = y$. Points along the upper-right edge of the set of feasible outcomes are Pareto-optimal or Pareto-efficient outcomes, i.e., making one bargainer better off entails making the other worse off. Point Q represents the status quo, which defines what the players receive if they cannot agree on a new allocation.

[3]Unlike noncooperative bargaining theory, which emphasizes the bargaining process, cooperative or axiomatic bargaining theory generally focuses on the properties of a bargaining outcome. In particular, this approach specifies a priori properties or axioms that agreements are assumed to satisfy and then looks for feasible divisions of the surplus that satisfy these conditions. For example, the Nash bargaining solution posits that the outcome will be Pareto-optimal, whereas the noncooperative approach specifies the bargaining setting and then asks whether this setting leads to Pareto-efficient outcomes.

In addition to specifying the stakes, the noncooperative approach also requires specifying a bargaining protocol—how the players bargain about these stakes. Three protocols are widely used in applied work. In the first, player *1* makes a take-it-or-leave-it offer. If *2* accepts, both players receive the agreed payoffs; if *2* rejects, the status quo remains in place. In the second protocol, only one bargainer makes offers, but now that bargainer can make as many offers as she wants. In the third, offers alternate back and forth. If *2* rejects an offer, she can then make a counteroffer. If *1* rejects the counteroffer, he can counter the counter, and so on. There is typically no limit to the number of offers.

In addition to describing who can make offers and in what order, the bargaining protocol also specifies the other actions that the bargainers can take. One possibility is especially relevant to the recent work on war. Sometimes one or both bargainers have an outside option that they can pursue after terminating the bargaining. A seller, for example, might stop bargaining with one potential buyer in order to start bargaining with another. One litigant might give up negotiating a settlement and go to trial. One state might stop bargaining and try to use force to impose a settlement.

In general, what happens if one player terminates the bargaining is not modeled explicitly (see Fudenberg et al. 1987 for an exception) and is simply abbreviated in the payoffs. That is, if one of the bargainers stops the bargaining, the game ends and the players receive the payoffs associated with the outside option. Point Ω denotes these payoffs in Figure 1. Note further that player *2* prefers the outside option Ω to the status quo Q whereas *1* prefers the status quo. As will be seen, these preferences affect the players' ability to make credible threats to exercise the outside option and thereby obtain a more favorable agreement.

The outside-option payoffs Ω are what the bargainers obtain if they fail to reach an agreement because they terminate the bargaining. By contrast, the players receive the payoffs associated with the status quo as long as they have failed to reach an agreement but have not yet ended the bargaining by pursuing an outside option. For this reason, the status quo is sometimes referred to as an inside option (Muthoo 1999, pp. 137–43).

How do different bargaining settings affect the outcome? Rubinstein (1982) studied a situation in which two players were trying to decide how to divide a "pie" and got nothing if they could not agree on the division. [This means $Q = (0, 0)$ in terms of Figure 1.] In his alternating-offer, infinite-horizon model, the players took turns making offers and there was no limit on the number of offers allowed.[4] The bargainers also had complete information about the bargaining setting, and, in particular, each knew the other's payoffs. Rubinstein proved two remarkable facts. First, although the game has infinitely many Nash equilibria, it has a

[4]Ståhl (1972) studied the less natural case in which the bargainers could only make a predetermined number of offers.

unique subgame perfect equilibrium.[5] Moreover, as the time between offers becomes arbitrarily small, the payoffs associated with this outcome converge to the Nash bargaining solution, which in this case is ($\frac{1}{2}$, $\frac{1}{2}$).[6]

The intuition for these results is straightforward. If offers alternate back and forth and can do so without limit, then in effect each player alternates between two roles. A player is either making an offer or receiving one, and the game always looks the same whenever a player assumes one of these roles. Let m and r, respectively, be the equilibrium payoffs to a player who is making an offer and to a player receiving an offer. If a player accepts an offer, he obtains r. If, by contrast, he rejects an offer, he assumes the role of the offerer. The payoff to this role is m, except that it must be discounted because time passes between the the player's rejection of an offer and his subsequent counteroffer. Let δm denote the discounted value of obtaining m after this delay, where δ is the players' common discount factor. Then, a player is choosing between r and δm when deciding whether to accept an offer. Knowing this, the offerer "buys" acceptance at the cheapest possible price by offering the lowest price the receiver would accept. This means that the offerer must give the receiver a payoff r that satisfies $r = \delta m$, which leaves the offerer with what is left, namely, $m = 1 - r$. Solving these two equations for the equilibrium payoffs m and r gives $m = 1/(1 + \delta)$ and $r = \delta/(1 + \delta)$.

Now suppose that the time before a bargainer can make a counteroffer is arbitrarily small. This means the receiver pays almost nothing to reject an offer and thereby become the offerer. Formally, δ becomes arbitrarily close to 1 as the time between offers becomes arbitrarily small, and as δ goes to 1, $(m, r) = (1/(1 + \delta), \delta/(1 + \delta))$ goes to ($\frac{1}{2}$, $\frac{1}{2}$). More substantively, as the time between offers becomes very small, there is virtually no difference between the roles of making and receiving an offer, for anyone in the latter role can always take on the former by rejecting the offer at little cost if there is a short time between rounds. Thus, the unique equilibrium gives identical players identical payoffs. These stunning results—uniqueness and convergence to the Nash bargaining solution—renewed interest in noncooperative bargaining theory and led to an explosion of work.

When offers alternate back and forth and the time between offers is small, the bargainers are in almost identical situations and therefore have about the same bargaining power. In these circumstances they divide the surplus or pie in half. When only one player makes all the offers, that player has all the bargaining power. Suppose player *1* can make a take-it-or-leave-it offer. If player *2* rejects it,

[5]Unlike a Nash equilibrium, a subgame perfect equilibrium requires the threats implicit in the bargainer's strategies to be credible. Because they exclude incredible threats and promises, subgame perfect equilibria offer more plausible predictions about outcomes than do Nash equilibria.

[6]Binmore (1987), Muthoo (1999), and Osborne & Rubinstein (1990) discuss the Nash bargaining solution, and Roth (1979) discusses axiomatic bargaining theory more generally.

he obtains zero. Exploiting this, *1* claims all of the surplus for herself by offering *2* zero.[7] The same result obtains if *1* can make more than one offer.

The existence of outside options can affect a player's bargaining power. Suppose that *1* makes a take-it-or-leave-it offer. In response, *2* can accept or reject the offer or exercise an outside option that ends the game and yields the payoffs associated with Ω in Figure 1. If *2* did not have this option, *1* would maximize her payoff by offering *2* the smallest share that he would be willing to accept and claiming the rest. That is, *1* would propose $P_1 = (1 - q_2, q_2)$. If, by contrast, *2* has the outside option Ω, which he prefers to the status quo (since $\omega_2 > q_2$), then he can credibly threaten to exercise the outside option if he is offered anything less than ω_2. Understanding this, *1* proposes $P_\Omega = (1 - \omega_2, \omega_2)$. In this case, the existence of an outside option and *2*'s ability to credibly threaten to exercise it improves his bargaining position and gets him better terms. The same outcome results if *1* can make more than a single offer.

The situation is different if, as in the Rubinstein model, offers alternate. Suppose that when considering an offer, a bargainer can accept it, reject it in order to make a counteroffer, or exercise the outside option Ω. The outcome of this game would be A in Figure 1 if the players did not have this option and if the interval between offers were very short. (Point A, the Nash bargaining solution relative to threat point Q, divides the bargaining surplus evenly between the bargainers relative to the status quo Q.) Note further that the bargainers prefer A to Ω. That is, both bargainers prefer the agreement they would reach absent an outside option to the payoffs associated with that option. In these circumstances, neither bargainer can credibly threaten to exercise the outside option, and that option has no effect on the bargaining outcome. A is the outcome regardless of the presence of Ω.[8] Matters are different if the outside option is Ω'. Now *2* prefers Ω' to A and therefore can credibly threaten to exercise the outside option unless offered ω_2', which is strictly greater than a_2. In this situation, *1* offers *2* just enough to make the exercise of the outside option incredible, namely $A' = (1 - \omega_2', \omega_2')$.

A striking feature of actual bargaining is that it often results in costly delays and inefficient outcomes. Haggling between a buyer and seller delays agreement. Labor negotiations break down in costly strikes. Litigants fail to reach out-of-court settlements and engage in expensive trials. States fall short in their diplomatic efforts to resolve a conflict and go to war. In all these cases, the outcome of the bargaining is not Pareto-optimal. Whatever the final agreement, both sides would have been better off agreeing to it at the very outset and thereby at least avoiding the bargaining costs. The eminent economist Hicks (1932) believed that these

[7]Player *2* clearly would accept any offer greater than zero. One can formally show that player *2* in equilibrium is sure to accept an offer of zero although there is no difference between accepting and rejecting it.

[8]Technically, this result also depends on the precise protocol and in particular on exactly when the bargainers can exercise the outside option (see Osborne & Rubinstein 1990, pp. 54–63).

inefficiencies resulted from irrational or misguided behavior. By the early 1980s, economists believed that incomplete or asymmetric information would provide a much better explanation, and this belief motivated a great deal of work.

The basic idea was that if, say, a seller was uncertain about how much a buyer was willing to pay for something, then he might begin by charging a high price and subsequently lowering it. Obviously a low-valuation buyer would not pay a high price and would wait for a lower one. But a high-valuation buyer might pay a higher price rather than wait for a lower one if the benefits of a buying at a lower price were outweighed by the costs of delaying an agreement. Indeed, if the buyer were sufficiently likely to agree to a high price, then it would be optimal for the seller to start the bargaining by demanding a high price and then gradually lower it. In this way, asymmetric information would explain delay.

Unfortunately, efforts to explain delay and other bargaining inefficiencies on the basis of asymmetric information have not been entirely successful. For example, as the time between offers becomes very small, bargainers generally reach agreement without delay even in the presence of uncertainty. Thus, it is not asymmetric information per se that accounts for delay, but the rather unsatisfying assumption that a significant amount of time must elapse before the seller can make a new offer. Asymmetric information is related to bargaining inefficiencies, but its limited ability to explain delay, which is the simplest kind of inefficiency, should serve as a note of caution that the work on war discussed below will have to address.[9]

WAR AS A BARGAINING PROCESS: THE BASIC FRAMEWORK

Much of the recent formal work on international conflict shares a common, unifying theme. The origin, conduct, and termination of war are part of a bargaining process. This perspective is, of course, not new. Schelling, perhaps most famously, observed that most conflicts "are essentially *bargaining* situations" (1960, p. 5). What is new is the set of game-theoretic tools that makes it possible to follow through on this perspective to a greater extent.[10] This section describes the bargaining-problem framework.

[9]Kennan & Wilson (1993) review the work on bargaining with private information. For an introduction to the problem of delay and the related issue of the Coase conjecture, see Fudenberg & Tirole (1991, pp. 397–434); Gul & Sonnenschein (1988); and Gul et al. (1986).

[10]To appreciate the importance of the new tools, note that uncertainty plays a crucial role in bargaining. Sellers, for example, are uncertain about buyers' unwillingness to pay; states are uncertain about each other's resolve. Despite its importance, no one knew how to study uncertainty and asymmetric information formally until Harsanyi's (1967–1968) work was combined with ideas about credibility and perfect equilibria in the early 1980s. These developments underpin the latest wave of work on bargaining theory as well as the explosion of work that revolutionized economics in the 1980s and 1990s. See Kreps (1990) for an accessible overview of these developments and some of their limitations.

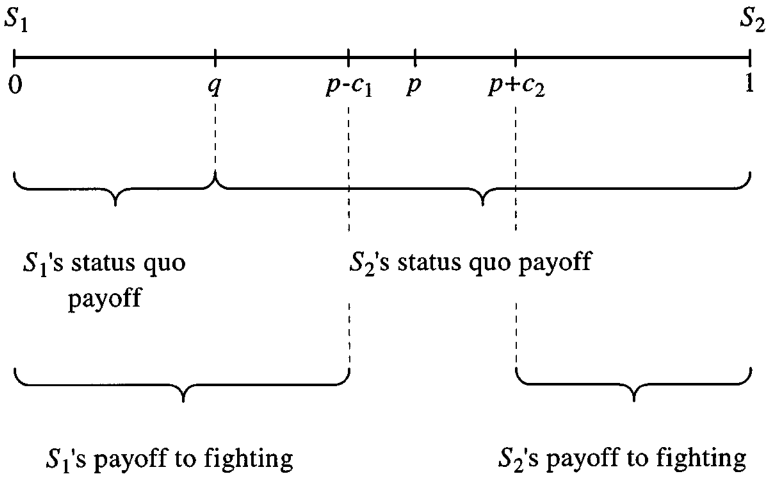

Figure 2 Bargaining over territory.

Figure 2 illustrates the basic setup. Two states, S_1 and S_2, are bargaining about revising the status quo. The bargaining can be about any issue, but it is usually taken to be about territory. In this interpretation, S_1 controls all territory to the left of q, from which it obtains utility q. S_2 controls all territory to the right of q, from which it derives utility $1 - q$. The interval [0, 1] defines the range of possible territorial agreements, and the states receive utilities $U_D(x) = x$ and $U_S(x) = 1 - x$ by agreeing to $x \in [0, 1]$. (Bargaining models typically assume states maximize their absolute gains. See below for a discussion of the implications of "relative-gains" concerns.)

In addition to revising the status quo through mutual agreement, the states may also use force to try to impose a settlement. If they fight, S_1 pays cost c_1 and wins all the territory with probability p. With probability $1 - p$, S_1 loses everything and also pays cost c_1. Thus, S_1's expected payoff to fighting is $p(1 - c_1) + (1 - p)(0 - c_1) = p - c_1$. Similarly, S_2's payoff to fighting is $1 - p - c_2$. In this setting, it is natural to interpret p as the distribution of power between S_1 and S_2.[11]

In Figure 2, S_1 prefers fighting to accepting any point to the right of $p - c_1$ and prefers accepting any point to the left of $p - c_1$ to fighting. Similarly, S_2 prefers the distribution y to fighting if $1 - y > 1 - p - c_2 \Leftrightarrow p + c_2 \leq y$. Consequently, S_1 is dissatisfied with the status quo, i.e., prefers fighting to accepting q, if $q < p - c_1$, whereas S_2 is satisfied since $q \leq p + c_2$. Thus, the set of feasible peaceful

[11]The assumption that one state or the other wins everything has no effect on the formal analysis, since the results are the same if p is taken to be the expected territorial outcome. However, if p is defined that way, it may no longer make sense to think of p as the distribution of power. Suppose, for example, the expected territorial outcome remains the same but the variance of the outcome goes up. Is the distribution of power the same or not?

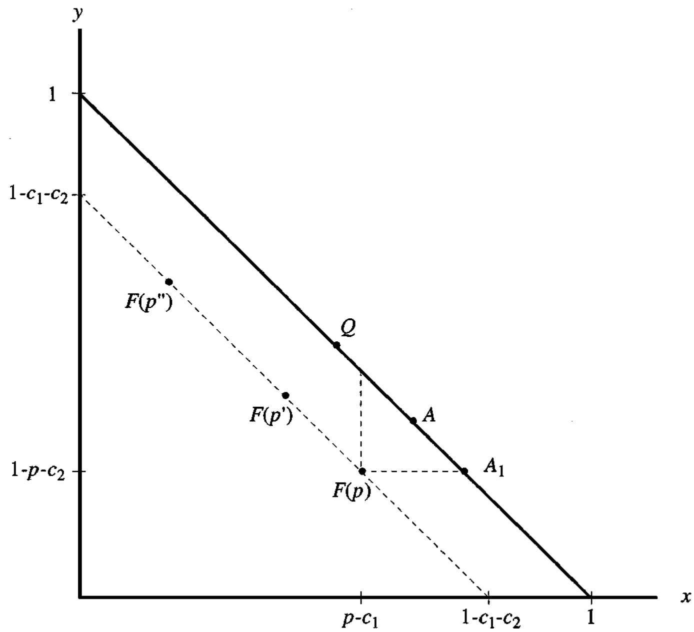

Figure 3 War as an outside option.

agreements, i.e., territorial divisions that both states prefer to fighting, lies between $p - c_1$ and $p + c_2$.

Figure 3 recasts the bargaining problem. S_1's and S_2's utilities are plotted along the horizontal and vertical axes, respectively. The set of peaceful outcomes, including the continuation of the status quo Q, are on the line between (1, 0) and (0, 1) and define the Pareto frontier of the bargaining problem. If the states fight, they obtain the payoffs at $F(p)$. That this outcome lies inside the Pareto frontier reflects the fact that fighting is costly and therefore inefficient. Nevertheless, S_1 prefers $F(p)$ to Q, since the former lies to the right of the latter. The allocations on the Pareto frontier above and to the right of $F(p)$ are the peaceful outcomes that both states prefer to fighting.

Figure 3 shows how shifts in the distribution of power affect the bargaining problem. As the distribution of power shifts in favor of S_2, say from p to p' to p'' (where $p > p' > p''$), S_2's payoff to fighting increases, S_1's decreases, and $F(p)$ slides upward along the line from $(1 - c_1 - c_2, 0)$ to $(0, 1 - c_1 - c_2)$.[12] At p', both states prefer Q to $F(p')$ and neither is dissatisfied. At p'', by contrast, S_2 prefers $F(p'')$ to Q, while S_1 prefers Q and is satisfied.

[12]If $p = c_1$, then S_1's and S_2's expected payoffs are 0 and $1 - p - c_2 = 1 - c_1 - c_2$, respectively. Thus, $F(c_1) = (0, 1 - c_1 - c_2)$, and similarly, $F(1 - c_2) = (1 - c_1 - c_2, 0)$.

In an important article, Fearon (1995) uses the basic bargaining setup in Figure 2 and a take-it-or-leave-it protocol to reframe the theories of the origin of war by linking them to a more general problem in bargaining theory, namely that of explaining why bargaining ever breaks down in inefficient outcomes. Because fighting is costly, $F(p)$ lies inside the Pareto frontier and there are agreements that *both* states prefer to fighting (e.g., A in Figure 3). A rationalist theory of war, Fearon argues, must explain why states end up at $F(p)$ and not at a peaceful settlement such as A, which makes them both better off.

Once the question is posed this clearly, it is immediately evident that three of the then most prominent theories of war fail to address this fundamental puzzle. Theories that appeal to anarchy as an important structural cause of war (e.g., Waltz 1959, 1979) assume that there is nothing to stop states from using force to further their ends if doing so appears to be in their best interest. But the puzzle is to explain why states use force when it is not in their interest, i.e., there are outcomes that both prefer to fighting. The idea that a state goes to war when it has a positive expected utility for fighting (e.g., Bueno de Mesquita 1981) also falters. S_1 has a positive expected utility for fighting, i.e., it prefers $F(p)$ to q, but this does not explain why the states would end up at $F(p)$ rather than A, since the expected utility of A is higher for both states than that of $F(p)$. Finally, the theory of preventive war argues that a declining state may attack a rising power in order to avoid having to fight later on worse terms (see Levy 1987 for a survey of nonformal theories of preventive war). This perspective introduces a dynamic component, namely the effects of shifts in the distribution of power on bargaining. But as it stands, this theory also fails to explain why bargaining breaks down in inefficient outcomes. For as long as fighting is costly, the "pie" is larger and there is more to be divided if the states avoid fighting. This means that there should be divisions that leave both sides better off.

In addition to these three theoretical schools, offense-defense theory and relative-gains concerns also suffer from the same weakness. As shown below, neither explains why states use force when there are Pareto-superior alternatives that the states prefer to fighting (see Lynn-Jones 1995 and Van Evera 1998 for reviews of offense-defense theory).

Fearon suggests that coherent rationalist explanations will take one of two general forms.[13] The first appeals to asymmetric information. Suppose S_1 can make a take-it-or-leave-it offer, which S_2 can accept or reject by fighting. If the states have complete information, then S_1 maximizes its payoff by making the largest demand that S_2 will accept. To wit, S_1 demands the border $p + c_2$ in Figure 2, which leaves S_2 with $1 - p - c_2$. S_2 cannot do better by fighting, so it accepts. As is typical in take-it-or-leave-it bargaining, the player making the offer obtains all of the surplus by leaving the other player indifferent between accepting and

[13]Fearon (1995) notes a third kind of rationalist explanation that is logically coherent but seems empirically unlikely. If the states are bargaining about an issue that is indivisible or "lumpy," then there may not be any feasible outcome that both states prefer to fighting.

rejecting the offer. If, by contrast, S_1 is unsure of S_2's cost of fighting c_2, S_1 no longer knows what it can demand of S_2 without provoking war. S_1 now faces a risk-return trade-off between possibly obtaining better terms and a higher probability of not obtaining any settlement at all. The more S_1 demands, the better off it will be if S_2 accepts. But the more it demands, the less likely S_2 is to accept. Typically, the optimal solution to this trade-off is not to "buy" zero risk. That is, the demand that maximizes S_1's expected payoff will be rejected with positive probability, in which case there will be war. Thus, war—or, more generally, inefficiency—results from asymmetric information.

The second explanation of why states may fail to agree on an outcome that both prefer to fighting is that they are unable to credibly commit themselves to following through on the agreement. The prisoner's dilemma is the classic example of this problem. Both actors prefer the cooperative outcome (C, C) to mutual defection (D, D), but this agreement is not self-enforcing, since at least one state, and in this case both, has an incentive to renege. Figure 3 makes the same point in a more general context. If none of the agreements above and to the right of $F(p)$ are self-enforcing, then the states may fight. Of course, the interesting part of this kind of explanation lies in explaining what about a particular strategic environment makes it impossible to credibly commit to these agreements. Fearon and others have addressed this in several settings.[14]

This very simple formulation can help cut through seemingly endless debates by posing central issues more clearly. Consider, for example, the idea that concerns about relative gains make cooperation more difficult and war more likely (Grieco 1988, Waltz 1979). Formally, S_1's and S_2's utilities for territorial division x are $U_1(x) = x - k(1 - x - x) = (1 + 2k)x - k$ and $U_2(x) = 1 - x - k[x - (1 - x)] = 1 + k - (1 + 2k)x$, where $k \geq 0$ measures the states' concern about relative gains.[15] Then S_1's payoff to fighting is $pU_D(1) + (1 - p)\,U_D(0) - c_1 = p(1 + 2k) - k - c_1$ and S_2's payoff is $1 - [(1 + 2k)p - k + c_2]$. Hence, the bargaining range, i.e., the set of territorial divisions that both states prefer to fighting, is given by $p(1 + 2k) - k - c_1 \leq x \leq (1 + 2k)p - k + c_2$. This implies that even if states are concerned about relative gains, there is still a set of agreements they prefer to fighting. Indeed, the length of this range, i.e., the difference between the upper and lower ends of this range, is just $c_1 + c_2$. This length does not depend on k. Nor does the probability

[14]It is important to stress that the rational choice approach (whether or not it is formalized mathematically) is based on the "methodological bet" (see Lake & Powell 1999) that trying to understand war as the outcome of instrumentally rational actors in a particular strategic setting will prove fruitful. This emphasis on the strategic setting contrasts with other "bets" that focus on other factors, such as psychological or cognitive factors (e.g., Jervis 1976), as did Hicks's (1932) explanation of strikes.

[15]As I have emphasized elsewhere (Powell 1994, p. 336; 1999, pp. 54–58), modeling relative-gains concerns through the utility function is really a reduced-form approach that begs the prior question of whether the international system actually induces relative-gains concerns. Many theorists claim it does, but this has yet to be shown deductively.

of war. More specifically, if S_2 is uncertain of S_1's cost c_1 and believes that this cost is uniformly distributed over $[\underline{c}_1, \overline{c}_1]$, and S_2 makes a take-it-or-leave-it offer, then the probability of war is $(\overline{c}_1 - c_2)/[2(\overline{c}_1 - \underline{c}_1)]$. Clearly, relative-gains concerns alone cannot explain why states fight rather than agree on peaceful divisions that both prefer to fighting.

Take-it-or-leave-it protocols are widely used because they do capture important aspects of bargaining and are easier to analyze. But they are also somewhat unnatural. The risk-return trade-off described above reflects the fact that the state receiving the offer can respond only by accepting it or going to war. In most circumstances a state could make a counteroffer or, at minimum, neither agree nor attack and simply wait. How would this more natural setting affect the bargaining?

Powell (1996a,b, 1999) addresses this issue in the context of a Rubinstein model with outside options. In that infinite-horizon, alternating-offer game, the state receiving an offer can accept it or reject it, as in the Rubinstein model, or end the game by exercising an outside option. This option is interpreted as trying to use force to impose a settlement and gives the players payoff $F(p)$. If one state rejects the other's offer and does not exercise its outside option of fighting, the round ends and that state makes the offer in the next round.

Powell studies the case in which there is two-sided incomplete information in that neither state knows the other's cost of fighting. Somewhat surprisingly, this game turns out to have a unique equilibrium,[16] in which the status quo remains unchanged and there is no risk of war if both states are satisfied. If one state is dissatisfied, then this state either accepts the satisfied state's initial offer or attacks. Anticipating this reaction, the satisfied state makes its optimal de facto take-it-or-leave-it offer. The probability of war is simply the probability that the dissatisfied state rejects this offer.

As Figure 3 suggests, both states are satisfied if the distribution of benefits mirrors the distribution of power. That is, if $p' = q$, then Q and F will lie on the same ray from the origin as do Q and $F(p')$. Consequently, Q is Pareto-superior to $F(p')$ and both states will be satisfied when the distribution of power is p'. By contrast, S_1 is dissatisfied at p when there is a large disparity between the distributions of power and benefits ($|p - q| \gg 0$), and S_2 is dissatisfied at p'', where there also is a large disparity between the distributions of power and benefits ($|p'' - q| \gg 0$).

This result undercuts many of the claims about the relationship between the distribution of power and the likelihood of war. The balance-of-power school (Claude 1962, Morgenthau 1967, Mearsheimer 1990, Wright 1965, Wolfers 1962) claims that war is least likely when there is an even distribution of power ($p = 1/2$), whereas the preponderance-of-power school (Blainey 1973, Organski 1968, Organski & Kugler 1980) argues that war is least likely when there is a preponderance of power ($p \approx 1$ or $p \approx 0$). (See Levy 1989 for a review of this debate and Wagner 1994 for a discussion of conceptual and modeling issues.)

[16]Bargaining games in which offers are made by a player with private information (e.g., one who knows her own cost of fighting while the other player does not) are typically plagued by a plethora of equilibria.

However, substantial empirical efforts to answer this question have yielded conflicting results. Siverson & Tennefoss (1984) find an even distribution of power to be more peaceful, as does Ferris (1973). By contrast, Kim (1991, 1992), Moul (1988), and Weede (1976) determine that a preponderance of power is more peaceful. Singer et al. (1972) find evidence for both claims depending on the historical period—an even distribution of power is more peace prone in the nineteenth century but less so in the twentieth. Mansfield (1992, 1994) uncovers evidence of a nonlinear, quadratic relationship in which the probability of war is smallest when there is both an even distribution of power and a preponderance of power. The greatest instability occurs somewhere between a preponderance and a balance of power. Finally, Maoz (1983) and Bueno de Mesquita & Lalman (1988) find no significant relation between stability and the distribution of power.

The formal results above may account for these conflicting results: The probability of war is likely to be related to the relationship between the distributions of power and benefits, not solely to the distribution of power. Consequently, any effort to assess the relationship between the distribution of power and the probability of war must control for the distribution of benefits.

This model also illuminates the effects of the offense-defense balance on the likelihood of war. Jervis' (1979) important and enormously influential article linked the security dilemma to the offense-defense balance and laid the foundation for what has become the offense-defense theory of war (see Glaser 1997 for a review of work on the security dilemma). The basic idea is that factors that make attacking relatively more attractive than defending make war more likely. Jervis framed his discussion in terms of a repeated prisoner's dilemma and a 2×2 stag hunt. But it is not immediately evident from a bargaining perspective why increasing a state's payoff to attacking should make war more likely. Suppose that the probability that S_1 prevails is $p + f$ if it attacks and $p - f$ if it is attacked. Then the difference between S_1's probability of prevailing if it attacks and its probability of prevailing if it is attacked is $2f$, so f can be thought of as the size of the offensive advantage. In terms of Figure 2, the presence of an offensive advantage narrows the bargaining range to the interval between $p + f - c_1$ and $p - f + c_2$. The length of this interval is $c_1 + c_2 - 2f$, which clearly decreases as the offense becomes more favorable. But as long as the size of the offensive advantage f is not too large, there will still be a set of agreements that both sides prefer to fighting. Why, then, do the states not agree to one of these? Offense-defense theory provides no general explanation. Powell's (1999, pp. 110–13) analysis of the asymmetric-information bargaining model shows that larger offensive advantages as well as decreases in the cost of fighting do make war more likely.

There is a natural link between bargaining and arms races: One can see a state's attempt to build up its military strength as an effort to create a more advantageous bargaining environment (Schelling 1966). Kydd (2000) pursues this line in his analysis of the "deterrence model" of arms racing. In the deterrence model (Jervis 1976), arms races are a symptom of an underlying conflict of interest between states. Consequently, the outbreak of war may be correlated with arms races, but the link is not causal. In the "spiral model" (Jervis 1976), by contrast, there is no fundamental conflict. Rather, misperceptions and arms races interact to fuel a spiral of hostilities that may eventually end in war. In the spiral model, arms races are causal in the sense that if one could stop the cycle of misperception, no state would want to attack the other. (Recent contributions to the empirical literature on arms racing include Diehl & Crescenzi 1998 and Sample 1997.)

In Kydd's (2000) formulation, a state, say S_1, makes a take-it-or-leave-it offer to S_2. If S_2 rejects, the states can choose to increase their military capabilities, after which there is another round of take-it-or-leave-it bargaining. If the states have complete information about each other's ability to sustain an arms race, there will be no arms racing because any mismatch between the distributions of power and benefits will be brought into line in the first round of bargaining. If, however, S_2 is unsure of S_1's ability to run an arms race, S_2 may reject S_1's initial demand because S_2 believes that S_1 is bluffing. In turn, S_1, if it is not bluffing, will build up its military capabilities in order to "signal" to S_2 that it is deadly serious. Kydd's analysis provides a firm, formal footing for the deterrence model. The challenge now is to integrate this work with that on the spiral model in an effort to assess their relative empirical importance.[17]

MULTILATERAL BARGAINING: BALANCING AND BANDWAGONING, EXTENDED DETERRENCE, AND THIRD-PARTY INTERVENTION

Most of noncooperative bargaining theory, as well as the preceding applications, focuses on two-party bargaining. This emphasis probably reflects the fact that although n-player games are not necessarily any more difficult to analyze formally, it is often harder to specify a substantively convincing or "natural" bargaining protocol. For example, it seems natural in a two-player, Rubinstein game for the players to take turns making offers (although, in fact, there is nothing really natural about this). But what is the "natural" bargaining protocol if there are three players? Should the offers be made round-robin, i.e., *1* makes an offer, then *2*, then *3*, then *1* again and so on? This hardly seems natural. Suppose, instead, that the player who makes the next offer is chosen randomly. That is, *1*, *2*, or *3* each make the first offer with probability $1/3$. The player making the next offer is then selected randomly again (see Baron & Ferejohn 1989 for an important analysis of parliamentary bargaining based on this kind of protocol). This protocol has the advantage of symmetry. But it does not feel very natural.

Nevertheless, recent work has begun to study interactions between three or more actors, sometimes finessing the problem by formalizing the situation so that

[17]See Kydd (1997) for a formal analysis of the spiral model. Because this analysis is not based on a bargaining model, it is not reviewed here.

bargaining only occurs between two actors. Two areas of research are discussed here. The first is central to balancing, bandwagoning, and states' alignment behavior. The second is a small but growing body of work on third-party intervention in military disputes.

The idea that states balance against power can be traced back to Thucydides and Xenophon and to the politics of ancient Greece—or, at least, David Hume (1898 [1752]) thought so. According to Mattingly (1955), balance-of-power politics framed the diplomacy of Northern Italy in the late fifteenth century.[18] More recently, Waltz observed, "If there is any distinctively political theory of international politics, balance-of-power theory is it. And yet one cannot find a single statement of it that is generally accepted" (1979, p. 117).

Wagner's (1986) path-breaking article renewed interest in trying to understand balancing and bandwagoning as the equilibrium outcome of an underlying noncooperative dynamic game.[19] He sketched a very complicated n-player game, which Niou & Ordeshook (1990) refined and analyzed more extensively. In their game, states bargain by proposing coalitions. A state is chosen randomly to make a proposal, and one coalition prevails over another if the resources of its members exceed the resources of the other's members. The primary finding is that no state is ever eliminated as long as it is an essential part of some winning coalition. This disappointing result suggests that balance-of-power politics has very few observational implications. Unfortunately, these models are so complex, as indeed they must be in order to do what Wagner and Niou & Ordeshook want them to do, that it is difficult to know what is driving the results.

Powell (1999, pp. 149–96) looks at a much simpler game that is more closely tied to the bilateral bargaining models discussed above. In his formulation, a potential attacker, A, can attack S_1 only, S_2 only, both S_1 and S_2, or neither. If A attacks only one of the other states, say S_2, then S_1 has the following options: (a) balance against A by joining S_2, (b) bandwagon with A by joining the attack on S_2, or (c) stand aside while A and S_2 fight. As in the models above, fighting always results in the elimination of one of the opposing sides. Thus, at most two states will remain after the first round of fighting. These two states then bargain about revising the territorial status quo given the distribution of power and benefits that results from the outcome of the first round of fighting.

This formulation highlights the trade-off a state faces when deciding whether to balance or bandwagon. Balancing with a weaker state puts the balancer in a stronger bargaining position relative to its coalition partner and thus makes for a more favorable division of the spoils of victory—if this coalition prevails against the third state. By contrast, aligning with the stronger state puts the bandwagoner in a weaker bargaining position relative to its coalition partner and makes for a less

[18]See Haas (1953), Hinsley (1963), and Knutsen (1997) for historical overviews of balance-of-power theories and thinking. Butterfield (1966, p. 139) offers a skeptical view of the existence of a coherent conception of balancing before the middle of the seventeenth century.
[19]See Kaplan et al. (1960) for a very early, partially game theoretic effort to study balancing.

favorable division of the spoils of victory. But bandwagoning, relative to balancing, increases the probability of being on the winning side and having any spoils to divide.[20] Balancing in Powell's formulation turns out to be much less likely than bandwagoning.

Werner's (2000) interesting analysis focuses on the attacker's decision. In her model, an attacker, *A*, makes "offers" by choosing the size of the stakes of a dispute it is thinking about provoking with another state. The latter is the protégé of a third state, *D*, and *D* must decide whether to intervene on its protégé's behalf if *A* attacks. Since the stakes are endogenous, *A* faces a risk-return trade-off. The more it demands, the more likely *D* is to intervene. This trade-off induces selection effects that help to explain some empirical anomalies regarding intervention. For example, if *D* is relatively powerful, *A* will moderate its demands and thereby reduce the chances that *D* intervenes. Consequently, there should be no strong relation between *D*'s strength and the likelihood that it will intervene. This runs counter to the intuitive conjecture that more powerful states are more likely to intervene (Altfeld & Bueno de Mesquita 1979; Walt 1988, 1992; Labs 1992) and helps account for the empirical finding that they are not (Huth & Russett 1984).

The role of third-party mediators in dispute resolution has recently received a good deal of empirical attention (see Bercovitch 1996, Dixon 1996, and Kleibor 1996 for recent efforts and reviews). Kydd[21] and Rauchhaus[22] analyze this problem formally in order to provide a firmer theoretical foundation for the empirics. In Kydd's formulation, a state, say S_1, can make a take-it-or-leave-it demand of S_2. The former is uncertain of the latter's cost of fighting and so faces a risk-return trade-off. Kydd introduces a mediator into this setup by allowing the mediator to report its beliefs about these costs to S_1 before S_1 makes its demand. The mediator is also uncertain of S_2's cost but does have some independent information it can pass on.

Kydd defines a mediator to be biased in favor of one of the states if it prefers territorial distributions that favor that state. The mediator is unbiased if it does not care about the terms of a territorial settlement and only wants to minimize the risk of war. Kydd then shows that, surprisingly, a mediator must be biased in order

[20]Awareness of this trade-off is of course not new and not the result of formal analysis. But, strangely, some widely accepted claims that states generally balance appear to disregard the trade-off. Waltz, for example, argues that secondary states balance because they are "both more appreciated and safer, *provided, of course, that the coalition they join achieves enough deterrent or defensive strength to dissuade adversaries from attacking*" (1979, p. 127, emphasis added).

[21]Kydd A. 2001. Which side are you on? Mediation as cheap talk. Unpublished manuscript, Department of Political Science, University of California, Riverside.

[22]Rauchhaus RW. 2000. *Third-party intervention in militarized disputes: primum non nocere*. Unpublished PhD dissertation, Department of Political Science, University of California, Berkeley.

to have any effect on the outcome of the bargaining. The intuition is that if the mediator only cares about preventing war, then the mediator will always tell S_1 that S_2 is resolute (regardless of whether the mediator believes it) and therefore S_1 should moderate its demands. This advice, if followed, minimizes the probability of war. Of course, S_1 understands that the mediator has an incentive to lie and consequently discounts its advice.

Rauchhaus (2000, see footnote 22) generalizes this result. Suppose that a mediator is so strongly biased in favor of a state, or, perhaps more accurately, a particular territorial outcome, that it is willing to act as an *agent provocateur*. During the Cold War, for example, the United States tried to exclude the Soviet Union from a mediating role out of concern that the Soviet Union would "stir up trouble." This type of "mediator," the opposite of the one who prefers a peaceful outcome to any other, would encourage S_1 to make maximal demands. Kydd excludes this case, whereas Rauchhaus allows for it by letting the motivations of the mediator range from caring only about avoiding war to caring only about the distribution of territory.[23] Rauchhaus shows that in order to sustain an equilibrium in which the mediator tells the truth, it must pay a "reputational" cost if it lies (and, presumably, will be caught out at some future time). Moreover, the reputational cost needed to sustain a truth-telling equilibrium is smallest when the mediator is unbiased. This suggests, contra Kydd, that unbiased mediators are preferable. Kydd's and Rauchhaus's formulations are not completely comparable and much work remains to be done. But they have opened up an interesting avenue of work on the increasingly important post–Cold War problem of mediation and intervention.

STRUCTURING APPEALS TO DOMESTIC POLITICS

One of the oldest debates in international relations theory is over the relative importance of domestic and structural explanations (see Fearon 1998b for the distinction between these types of explanation). Definitions of what counts as a domestic or structural explanation vary, but at a minimum, structural theories treat states as unitary actors. However defined, domestic and structural explanations are typically thought to be rivals. But this is misguided. Although some structural theories (e.g., Waltz 1979) seem to suggest that one can explain at least the outline of state behavior without reference to states' goals or preferences (except possibly the very general goal of survival), this assertion runs counter to most of the recent formal work in international relations theory (as well as much of the nonformal work).[24] In order to specify or close a game theoretic model, the actor's preferences and beliefs must be defined. Moreover, most conclusions derived from these models turn out to be at least somewhat sensitive to the actor's preferences and, especially,

[23]Kydd (2001, p. 37, see footnote 21) implicitly excludes the latter possibility when he assumes that if S_1 prefers a given settlement to the risk of war, then so does the mediator.
[24]Powell (2002) addresses this notion of structural explanation.

beliefs. This dependence suggests a more fruitful way to think about the relation between domestic and structural explanations. Domestic politics is terribly complicated, and it is not at all clear what aspects of it are likely to have significant international effects. Treating states as unitary actors creates a baseline and helps isolate domestic factors that are most likely to provide significant explanatory leverage.

Schultz's (1998, 2001) work on the democratic peace is an exciting example of the fruitfulness of this perspective. The empirical finding that democratic states do not fight each other is quite robust (see Schultz 2001 for a review and Gowa 1999 for the most powerful challenge to the democratic peace thesis). But we do not yet have an empirically established, theoretical explanation of the democratic peace, i.e., an explanation that predicts other empirical patterns that have been verified. Schultz's effort to provide one grows out of the work surveyed above that identifies asymmetric information as a critical cause of war. If war results from asymmetric information and if there are systematic differences in the likelihood of war between democratic versus nondemocratic states, he conjectures, then perhaps democratic institutions tend to moderate informational asymmetries between states during a time a crisis. With that in mind, Schultz develops a simple model of crisis bargaining in which there is an opposition party in a democratic state. This party's preferences differ from those of the party in power in that, if nothing else, it would prefer to be in power. This difference in preferences (along with a relatively open press) makes it more difficult for the party in power in a democracy to make threats it is unwilling to carry out. Since it is the uncertainty regarding states' willingness to follow through that leads to breakdown and war, democracies should be less likely to be involved in war.

Schultz (1998, 2001) tests his informational explanation against a competing explanation of the democratic peace that is based on the notion that the costs of going to war are systematically higher for democratic states than for nondemocratic states. He shows that these two mechanisms make different predictions about a state's response to a challenge made by a democratic or a nondemocratic state. The informational mechanism predicts that a state is less likely to resist a challenge coming from a democratic state (because there is less uncertainty and the challenge is less likely to be a bluff) than one from a nondemocratic state, whereas the cost-based explanation predicts the opposite. Schultz finds strong empirical support for the informational mechanism.

Bueno de Mesquita et al. (1999a,b)[25] also break down the unitary-actor assumption in a fruitful way. They describe a domestic regime in terms of two dimensions. The first is the size of the "selectorate," which is the group that participates in the selection of a state's leader, and the second is the size of the winning coalition, which is the group whose support a leader must retain in order to remain in power. A democracy, for example, has both a large selectorate and a large winning coalition.

[25]See also: Bueno de Mesquita B, Smith A, Siverson RM, Morrow JD. 2001. *Staying Alive: The Logic of Political Survival*. Unpublished manuscript, The Hoover Institution, Stanford University.

In a monarchy or junta, both are small. An autocracy may have a large or small selectorate, but the winning coalition is always small.

Bueno de Mesquita et al. study how these different institutional settings affect the bargaining surrounding a leader's efforts to sustain a winning coalition. Suppose a leader has limited resources that he can allocate to the production of private or public goods. Private goods channel benefits to specific individuals, such as those inside the leader's coalition, whereas public goods benefit everyone. The leader wants to spend the minimal amount needed to remain in office and pocket the rest, and the bargaining proceeds accordingly.

The leader faces a trade-off. On the one hand, offering private goods is the best way of buying loyalty, because it creates a wedge between the benefits of those inside the coalition, who receive the private benefits, and those outside the coalition. This wedge increases the cost to defecting from the winning coalition and supporting someone else's bid for leadership. Buying political support through public goods does not create a wedge because both those inside and outside the winning coalition receive the benefits. On the other hand, providing benefits to a large number of people through private payoffs may be much more expensive than providing them with the same level of benefits through public goods. This suggests that a leader tends to maintain support through the provision of private benefits when the winning coalition is small, as in juntas or authoritarian regimes, and through public goods in democracies.

Bueno de Mesquita et al. (1999a; also see footnote 25) use this framework to explain the democratic peace. When confronted by international conflict, democracies resolve this trade-off by allocating more resources to fighting than autocracies do. This makes democracies unattractive targets and more selective in the states they threaten, thereby reducing the chances that democratic states engage in war.

WAR AS AN INSIDE OPTION

Most of the bargaining theory literature on the causes of war, including the work discussed above, formalizes war as an outside option in a bargaining game. Going to war is typically modeled as a game-ending move, the payoffs of which reflect the distribution of power and the states' costs of fighting. For instance, S_1's expected payoff to fighting in the example above was given by the costly lottery $p \cdot 1 + (1-p) \cdot 0 - c_1 = p - c_1$. Representing war as a costly lottery raises three issues that recent work is beginning to address by treating war as an inside option.

The first issue is whether modeling war in this way leads to misleading conclusions. All models make simplifying assumptions and are designed to answer some questions and not others. A simplification is neutral with respect to a set of questions if relaxing that assumption would not significantly affect the model's answers to those questions. By contrast, a simplification is distorting if relaxing it would significantly affect those answers. Modeling war as a costly lottery clearly simplifies the analysis in that it assumes away any further strategic interaction after the states go to war. To the extent that the anticipation of that interaction influences the states' pre-war behavior, failing to model intra-war interactions explicitly may lead to misleading conclusions about the causes of the initial decision to fight. In these circumstances, treating war as a costly lottery would be a distorting simplification.

Second, even if the costly-lottery assumption is neutral, this simplification makes it impossible to ask important questions about the strategic dynamics of inter-war behavior and war termination. This alone is a good reason for relaxing the assumption by modeling war as a costly process during which strategic interaction continues (see Wittman 1979 for a path-breaking effort to model conflict and war as a process).

The third issue raised by the costly-lottery assumption is more general. One of the advantages of casting the problem of the origins of war in terms of a bargaining breakdown is that this conceptualization, as noted above, links it to a number of other substantive and theoretical literatures. Ideally, the work on bargaining and war should help us understand the exercise of coercive power—be it economic, military, or political. But in order for the work on the causes of war to contribute more fully to our understanding of these other forms of coercion, and vice versa, it is important to relax the assumption that the imposition of the costly sanction of going to war is a game-ending move. Even if this assumption is a plausible first approximation for some analyses of the causes of at least major war, it is much less plausible in other contexts, where the issue in dispute (e.g., trade policies) is only one of many issues over which the bargainers continue to interact while applying costly coercive pressure.

Wagner (2000) challenges the game-ending, costly-lottery assumption in two ways. First, following Blainey (1973), Wagner argues that most wars arise because of uncertainty over the distribution of power and continue until the belligerents' perceptions of the distribution of power come into line. However, most existing formal models of bargaining and war (e.g., Fearon 1995; Powell 1996a, 1999) appeal to uncertainty over states' costs or, more generally, their preferences or levels of resolve. It is not clear to what extent conclusions inferred from asymmetric information about preferences carry over to settings in which there is information about the distribution of power.

Wagner's second, broader concern begins with the observation that wars generally end because the states agree to stop fighting and not because the states are incapable of continuing to fight. "Thus to explain why wars occur one must explain why states must fight before reaching an agreement" (2000, p. 469). Wagner maintains that the costly-lottery assumption is distorting even with respect to questions about the origins of war and "can only lead to misleading conclusions" (2000, p. 469).

Smith & Stam,[26] Filson & Werner,[27] and Powell[28] unpack various aspects of the costly-lottery formulation in their efforts to model war as a costly process

[26]Smith A, Stam A. 2001. Bargaining through conflict. Unpublished manuscript, Department of Political Science, Yale University.
[27]Filson D, Werner S. 2001. Bargaining and fighting. Unpublished manuscript, Department of Political Science, Emory University.
[28]Powell R. 2001. Bargaining while fighting. Unpublished manuscript, Department of Political Science, University of California, Berkeley.

during which the states can still bargain. Smith & Stam treat the distribution of capabilities in a much more sophisticated way than does the standard costly-lottery formulation. They assume that two states, S_1 and S_2, are competing for a prize and that there are N forts of which S_1 controls n and S_2 controls $N - n$ at the start of the game. At the start of each round, S_1 proposes a division of the prize, say x for itself and $1 - x$ for S_2. S_2 then accepts or rejects. Acceptance ends the game with the agreed division. Rejection ends the round in a battle in which S_1 wins one fort from S_2 with probability p and loses one fort to S_2 with probability $1 - p$, and both states pay a cost of fighting. Given this new division of forts, S_1 starts the next round by making another proposal, and the bargaining continues in this way until the states agree on a division or until one state loses all of its forts. The states also start the game with different beliefs about the value of p, which they update as the fighting continues.

Although still relatively simple, this way of formalizing the distribution of capabilities has a number of appealing properties (see Smith 1998 for an earlier, related formulation). First, the distribution of power, which may be taken to be the probability that S_1 captures all the forts given the number of forts it currently controls, can rise and fall instead of only changing monotonically as in Powell (2001, see footnote 28). One need only think of the initial German successes in World War II to see why the distribution of power should be able to rise and fall in a model. A second appealing property is that gains can accumulate: The more forts a state captures, the higher the probability that it will defeat the other state. The model may therefore shed light on the debate about the importance of accumulated and relative gains.

This formulation has many potential applications. Smith & Stam use it to study the likelihood of future conflict given a previous conflict. In their model, the longer any previous conflict lasted, the less likely future disputes are to end in war. As in the standard costly-lottery formulations, conflicts arise because of different beliefs about the expected payoffs to fighting. These beliefs converge over the course of the conflict, and the longer that conflict lasts, the closer together both states' beliefs will be. If, however, fighting is costly, a conflict may not last very long and beliefs may still be far apart when it ends. This disparity between the states' beliefs is the seed for future conflict.

Unfortunately, this analysis treats beliefs in a nonstandard game theoretic way, as the authors acknowledge (Smith & Stam 2001, p. 9, see footnote 26). To wit, each state knows that the other has different beliefs about p; indeed, it knows precisely what those beliefs are. Yet, each state disregards this knowledge when updating its own beliefs about p.[29] Filson & Werner (2001, see footnote 27) open up the costly-lottery assumption in a different way. For Smith & Stam,

[29]More formally, the states do not share a common prior over p. A discussion of the common-prior assumption in game theory is beyond the scope of this review. Suffice it to say that this assumption is extremely strong and, in my view, quite problematic. But disregarding it also raises serious consistency issues.

the distribution of power shifts as the distribution of forts changes, but fighting does not destroy the forts. For Filson & Werner, the states start out with limited military resources that are destroyed by fighting, and a state is eliminated when all of its military resources have been destroyed. S_1 begins by making an offer to S_2, which can either reject the offer or accept it and thereby end the game. If S_2 rejects, S_1 can end the game by quitting or continue by attacking. If the latter, then S_1 wins a battle with probability p and S_2 wins with probability $1 - p$. Winning and losing both destroy resources but differentially so. The bargaining and fighting continue in this way until the states reach agreement or one of them is eliminated. There is also one-sided asymmetric information about the distribution of power. S_1 does not know p.

Filson & Werner (2001, see footnote 27) analyze this model in two kinds of cases. In the first, S_1 cannot sustain any losses and is eliminated as soon as it loses a battle; S_2 is eliminated after losing two battles; and S_1 is uncertain of p (i.e., of S_2's type) but believes that it is either one of two values. These restrictions mean that the states can fight no more than two battles before the game ends. The second kind of case is a set of numerical examples. In particular, the authors posit the number of types, the specific values of p associated with each type, and S_1's distribution over these types.

Beliefs are treated in the standard way with S_1 using all of its information to update its beliefs. As the fighting progresses, victories make S_1 more confident that it is facing weaker (lower p) types, whereas defeats make it more confident that it is facing tougher types. Consequently, the distribution of power—that is, the probability that a state will be eliminated—can rise and fall, and gains can accumulate as in Smith & Stam (2001, see footnote 26). This tends to make S_1 more likely to propose an offer that is sure to be accepted by all types following a defeat.

Powell (2001, see footnote 28) opens up the costly-lottery assumption in yet a third way. He studies an alternating-offer, infinite-horizon model. At the start of each round, state S_1 makes a proposal to which S_2 can respond by accepting the offer, waiting, or fighting. Accepting the offer ends the game in the proposed division of benefits. If S_2 fights, S_1 and S_2 pay costs c_1 and c_2, respectively. Fighting also generates some possibly very small risks (k_1 and k_2) that S_1 and S_2, respectively, collapse militarily. If one state collapses and the other does not, the latter prevails and obtains all the benefits. If both collapse simultaneously, the status quo division of benefits remains unchanged. If neither state collapses, the round ends and the next begins with another offer from S_1. If S_2 waits in response to S_1's offer, S_1 can fight or wait. The consequences of fighting are as before: The states pay costs c_1 and c_2 and generate risks of collapse k_1 and k_2. If S_1 waits, the round ends and the next begins with S_1's making another offer.

Powell analyzes the cases in which S_1 is uncertain about S_2's cost of fighting (c_2) and in which S_1 is unsure about S_2's probability of collapse (k_2). If one defines the distribution of power as the probability that a state prevails in a fight to the finish, then the distribution of power is a function of the probabilities that the states collapse. Hence, S_1's uncertainty about k_2 is equivalent to being uncertain about the distribution of power.

The equilibrium dynamics in these cases turn out to be similar. If the distribution of benefits mirrors the distribution of power, neither state prefers fighting to accepting the status quo, and the status quo remains unchanged. If, however, S_2 is dissatisfied, then in equilibrium S_1 makes a series of strictly increasing concessions. These offers "screen" S_2 according to its type. Weaker types, i.e., those with higher costs or larger probabilities of collapse, accept earlier offers. Tougher types never pass and always fight until they eventually accept a more favorable offer, albeit at the price of having had to fight longer.

There is, however, an interesting difference between these mechanisms. Suppose, in keeping with the economics literature, we examine what happens as the time between offers becomes small. The substantive idea here is that states can make and respond to offers much more quickly than they can prepare for and fight battles. Formally, S_1 can make n offers between battles while the states prepare to fight. Preparation is costly, with S_1 and S_2 paying c_1' and c_2' during each round (where these costs are defined so that the total cost of preparing for and fighting one battle remains c_1 or c_2 regardless of the number of offers between battles). However, the states fight and therefore only generate a risk of collapse in rounds $n, 2n, \ldots$. Paralleling the results obtained in buyer-seller models (e.g., Fudenberg et al. 1985, Gul et al. 1986), as the time between offers becomes small, the states reach agreement almost immediately and without fighting if the states are uncertain about costs. By contrast, if the states are uncertain about the distribution of power and if there is any delay, then there will be some fighting.

In sum, relaxing or unpacking the costly-lottery assumption is a natural next step in the development of the literature on bargaining and war, and much work is under way. This work promises to deepen our understanding of intra-war bargaining, war termination, and, ideally, the dynamics of coercive bargaining more generally. It is too soon to tell, but it also appears that the costly-lottery assumption is generally neutral with respect to questions about the origins of war and serves as a useful analytic simplification in some settings.

WAR WITH COMPLETE INFORMATION

Most of the recent work on bargaining and war has focused on the role of asymmetric information. But some work in comparative as well as international politics has begun to focus on commitment problems (Fearon 1995), especially dynamic commitment problems. As long as fighting destroys valuable resources, there is more to be divided between the bargainers if they can avoid fighting. Because there is more to go around if the states avoid fighting, there are divisions of this larger "pie" that are Pareto-superior to the expected outcome of fighting. Indeed, even in the presence of asymmetric information, there are generally divisions that are known to be Pareto-superior to fighting. Suppose, for example, S_2 makes all the offers to S_1 and is uncertain of S_1's cost of fighting but believes it is at least $\underline{c}_1$. Then S_2 can buy zero risk of war by appeasing the toughest possible type by offering $x = p - \underline{c}_1$. However, this offer is more than would be needed to appease higher-cost types, and asymmetric-information models explain war as the result of a risk-return trade-off in which the bargainers make smaller concessions at the cost of accepting some risk of breakdown.

This simple formulation provides a dubious reading of some important historical cases. If there were no uncertainty, then there would never be any war. Put another way, no matter how expansive an adversary's demands, a state would always prefer satisfying those demands to fighting as long as that state were sure of what the demands were and of precisely what it would take to appease its adversary. In symbols, S_2's payoff to appeasing an adversary with known cost c_1 by offering $x = p - c_1$ is $1 - x = 1 - p + c_1$, which is greater than its cost to fighting $1 - p - c_2$ as long as fighting is costly ($c_1 + c_2 > 0$). Uncertainty is surely an important cause of many wars, but a much more plausible reading of, say, the 1930s in Europe is that over time Britain and France became more confident of Germany's or Hitler's "type" and that this was a type they preferred to fight rather than appease. Many of the existing bargaining models miss this dynamic, because they assume that a state, if sure of its adversary's type, always prefers to appease it.

Dynamic commitment problems do not appeal to asymmetric information in order to explain bargaining breakdowns. Rather, they explain inefficiency and breakdown by the inability of the bargainers to abide by the Pareto-superior divisions because at least one of them will have an incentive to renege on or "renegotiate" any agreement. Shifts in the distribution of power or, more generally, in the cost of fighting are often at the heart of this dynamic-commitment approach or mechanism. When the distribution of power or the costs of fighting shift very quickly, then the larger pie is still not big enough to satisfy their minimal demands.

More formally, suppose that S_1 and S_2 are bargaining about a flow of benefits or series of pies. That is, a new pie has to be divided in each period. The value of this series (to risk-neutral bargainers) is $1 + \delta + \delta^2 + \cdots = 1/(1 - \delta)$, where δ is the states' common discount factor. Suppose further that if the states fight, S_1 wins the entire series with probability p and pays total cost $c_1/(1 - \delta)$, and S_2 wins everything with probability $1 - p$ at cost $c_2/(1 - \delta)$. In effect, S_1 can "lock in" the payoff or share $(p - c_1)/(1 - \delta)$ by attacking. That is, S_1 is indifferent between fighting and having the share $(p - c_1)/(1 - \delta)$ of the series. Hence, S_1 would never agree to a division that gave it a smaller payoff. Similarly, S_2 can "lock in" $(1 - p - c_2)/(1 - \delta)$ by fighting. The difference between these lock-ins, $(c_1 + c_2)/(1 - \delta)$, is what the states save by not fighting and is what they are bargaining about.

If S_2 makes all the offers and the distribution of power is expected to remain constant at p, then S_2 concedes just enough to make S_1 indifferent between fighting and accepting the offer. In symbols, S_2 proposes a division $x = p - c_1$ of each period's pie. But now suppose that the distribution of power will shift in the next period in S_1's favor so that the probability that S_1 would prevail will be $p + \Delta p$. If S_2 fights before this shift, it can still lock in $(1 - p - c_2)/(1 - \delta)$. Similarly, if S_1

waits one period until it is stronger and then attacks, it can lock in $(p - c_1) + \delta\,(p + \Delta p - c_1)/(1 - \delta)$, where the first term assumes S_2 offers $x = p - c_1$ to S_1 in the first period in order to prevent it from attacking then.

The sum of these lock-ins, $\delta(\Delta p - c_1 - c_2)/(1 - \delta)$, exceeds the total value there is to be divided, $1/(1 - \delta)$, if the distribution of power shifts more than the sum of the per-period costs of fighting, i.e., if $\delta\Delta p > c_1 + c_2$ or, more simply, $\Delta p > c_1 + c_2$. In these circumstances, bargaining breaks down in fighting even though there is complete information, i.e., each state knows the other's cost of fighting, because the larger pie created by not fighting is not large enough to be divided in a way that always gives each state at least as much *from that time forward* as the payoff it can lock in by fighting at that time. Because fighting destroys resources and there is less to be divided if the states fight, there are divisions that both sides would prefer to fighting if the states were sure to abide by those divisions in each and every period. But those divisions would at some point in the series specify an allocation that would give one state less than it could secure by fighting. Since the states cannot credibly commit themselves to not exploiting these situations, the bargainer that loses by waiting fights immediately.

Fearon (1998a) looks at ethnic conflict from the perspective of this mechanism. Imagine that a majority group M and a minority group m are bargaining about the extent of minority rights. Agreements are thought of as points on the interval [0, 1]; m prefers agreements closer to one and M prefers agreements closer to zero. If the groups fight, m prevails with probability p at cost c_m and M prevails with probability $1 - p$ at cost c_M. To keep the bargaining simple, assume that M can make a take-it-or-leave-it offer to m. If so, then M offers just enough to keep m indifferent between fighting and accepting, $x = p - c_m$.

Suppose, however, that after any agreement, majority groups generally consolidate their positions in a state. This effectively increases M's power and reduces m's to, say, $p - \Delta p$. This in turn reduces m's payoff to fighting, and this reduction gives M an incentive to renege on or renegotiate the agreement by offering $x' = p - \Delta p - c_m$. If we assume that M has been unable to commit itself to abiding by the original agreement, the choice facing m is a payoff of $p - c_m$, which it secures by fighting at the outset, or the lower payoff of $p - \Delta p - c_m$, which is what it ultimately would obtain through a peaceful and subsequently renegotiated settlement. Fighting is clearly better, so war results despite complete information. More substantively, this analysis very clearly highlights the importance of the majority group's ability to commit itself either through domestic or international institutions or guarantors. It also suggests a trade-off. The more the distribution of power is expected to change, the more credible these guarantees have to be (see Walter 1997 for an empirical assessment of this analysis).

At least since Thucydides, shifts in the distribution of power have been considered a significant source of international stress and potential conflict. Powell (1999) uses this basic dynamic-commitment approach to study this problem. As before, states are bargaining about revising the territorial status quo, but now the distribution of power is shifting in one state's favor throughout the bargaining. If, as in the example above, the distribution of power shifts quickly ($\Delta p > c_1 + c_2$), then the sum of the shares the states can lock in by fighting exceeds the amount there is to be divided. Bargaining breaks down in war in these circumstances despite complete information. If, by contrast, the distribution of power shifts more slowly, then the declining state makes a series of concessions to the rising state, and the bargaining does not break down unless there is also asymmetric information.

The same mechanism is at the heart of Acemoglu & Robinson's (2001) arguments about extending the franchise and about democratic transitions. Why, they ask, would an elite ever transfer political power to the masses in order to "buy off" social unrest? Why not buy them off through direct transfers financed by higher taxes on the elite—a more economically efficient policy that would not reduce the elite's political power? (As before, fighting, whether instigated by the elite or the masses, destroys resources, so there are agreements that all parties prefer to fighting.)

Acemoglu & Robinson (2001) argue that the transfer of political power serves as a means for the elite to commit itself to following through on its promise. That is, economic and social circumstances that create "revolutionary moments," which are formalized as times when the relative cost of challenging the elite is low, come and go. When a revolutionary moment passes and the cost of challenging the elite rises, there is nothing to keep the elite from reneging on its promises by ending the transfers. The promise of greater transfers is therefore incredible and cannot "buy off" pending social unrest and the threat of revolution. One way to make it credible is for the elite to give the beneficiaries of the promise the power to enforce it, which is what extending the franchise does. Of course, the leaders in a democracy may face the same kind of credibility problem in buying off potential coups, and Acemoglu & Robinson (2001) put this problem at the center of their analysis of democratic transitions. When breakdown does occur in these analyses, either through civil unrest or a coup, it is not the result of asymmetric information but of rapidly shifting lock-ins.

Finally, Fearon[30] looks at a different aspect of the way that shifting power distributions affect bargaining. In all the examples above, the reasons for the shift are unrelated to the bargaining settlement. Fearon examines the case in which shifts in the distribution of power arise endogenously because a gain in one period makes a bargainer stronger in the next. This prospect seems likely to create the kind of commitment problem illustrated above. A bargainer would fight today rather than make a concession because the concession makes it weaker tomorrow and more will be demanded of it. Surprisingly, this is not the case. Bargaining does not break down. Fighting today imposes immediate costs, whereas having to make concessions in the distant future is not very costly because of discounting. Thus, a state is willing to make concessions very slowly even if it ultimately

[30]Fearon JD. 1996. Bargaining over objects that influence future bargaining power. Unpublished manuscript, Department of Political Science, Stanford University.

has to concede a great deal (over the very long run), and it turns out that the other state also prefers this stream of concessions to bearing the immediate cost of fighting.

CONCLUSION

An apparently remarkable thing has happened in the past decade and a half. Twenty years ago Waltz lamented, "nothing in international politics seems to cumulate, not even criticism" (1979, p. 18). The work on bargaining and war is now a coherent, cumulating literature. Key puzzles have been framed as specific, well-defined problems. Generally, these problems have been studied first in the context of simpler models, which are subsequently refined and generalized and which raise previously unappreciated issues. These new issues are the springboard for new analyses. The literature on war and bargaining has taken on a self-sustaining quality that deals with issues spanning both international and comparative politics. So far, this work has been largely theoretical, as one would expect when new tools are first brought to bear. This theoretical work will continue and, ideally, be more fully complemented by serious empirical testing in an evolving modeling dialogue (Myerson 1992).

ACKNOWLEDGMENTS

Robert Jervis provided many helpful comments and criticisms, and I gratefully acknowledge the financial support of the NSF (SES-9911075).

LITERATURE CITED

Acemoglu D, Robinson JA. 2000. Why did the West extend the franchise? *Q. J. Econ.* 115:1167–99

Acemoglu D, Robinson J. 2001. A theory of political transitions. *Am. Econ. Rev.* 91:938–63

Altfeld MF, Bueno de Mesquita B. 1979. Choosing sides in war. *Int. Stud. Q.* 23:87–112

Baron DP, Ferejohn JA. 1989. Bargaining in legislatures. *Am. Polit. Sci. Rev.* 83(4):1181–1206

Bercovitch J. 1996. *Resolving International Conflicts*. Boulder, CO: Lynne Reinner

Binmore K. 1987. Nash bargaining theory. In *The Economics of Bargaining*, ed. K Binmore, P Dasgupta, pp. 27–46. New York: Blackwell

Blainey G. 1973. *The Causes of War*. New York: Free

Bueno de Mesquita B. 1981. *The War Trap*. New Haven, CT: Yale Univ. Press

Bueno de Mesquita B, Lalman D. 1988. Systemic and dyadic explanations of war. *World Polit.* 41:1–20

Bueno de Mesquita B, Morrow JD, Siverson RM, Smith A. 1999a. An institutional explanation of the democratic peace. *Am. Polit. Sci. Rev.* 93:791–807

Bueno de Mesquita B, Morrow JD, Siverson

RM, Smith A. 1999b. Political competition and economic growth. *J. Confl. Resolut.* 43:147–61

Butterfield H. 1950. The tragic element in modern international conflict. *Rev. Polit.* 12:147–64

Butterfield H. 1966. The balance of power. In *Diplomatic Investigations*, ed. H Butterfield, M Wright. London: Allen & Unwin

Claude I. 1962. *Power and International Relations*. New York: Random House

Crescenzi M, Diehl P. 1998. Reconfiguring the arms-race debate. *J. Peace Res.* 35:111–18

Dixon WJ. 1996. Third-party techniques for preventing conflict escalation and promoting peaceful settlements. *Int. Organ.* 50:653–81

Downs G, Rocke D. 1994. Conflict, agency, and gambling for resurrection. *Am. J. Polit. Sci.* 38:362–80

Fearon JD. 1994. Signaling versus the balance of power and interests. *J. Confl. Resolut.* 38:236–69

Fearon JD. 1995. Rationalist explanations for war. *Int. Organ.* 49:379–415

Fearon JD. 1998a. Commitment problems and the spread of ethnic conflict. In *The International Spread of Ethnic Conflict*, ed. DA Lake, D Rothchild, pp. 107–26. Princeton, NJ: Princeton Univ. Press

Fearon JD. 1998b. Domestic politics, foreign policy, and theories of international relations. *Annu. Rev. Polit. Sci.* 1:289–313

Ferris W. 1973. *The Power Capabilities of Nations*. Lexington, MA: Lexington Books

Fudenberg D, Levine D, Tirole J. 1985. Infinite-horizon models of bargaining with one-sided incomplete information. In *Game Theoretic Models of Bargaining*, ed. A Roth, pp. 73–98. New York: Cambridge Univ. Press

Fudenberg D, Levine D, Tirole J. 1987. Incomplete information bargaining with outside opportunities. *Q. J. Econ.* 102:37–50

Fudenberg D, Tirole J. 1991. *Game Theory*. Cambridge, MA: MIT Press

Glaser CL. 1997. The security dilemma revisited. *World Polit.* 50:171–210

Gowa J. 1999. *Ballots and Bullets*. Princeton, NJ: Princeton Univ. Press

Grieco J. 1988. Anarchy and the limits of cooperation. *Int. Organ.* 42:485–507

Gul F, Sonnenschein H. 1988. On delay in bargaining with one-sided uncertainty. *Econometrica* 56:601–11

Gul F, Sonnenschein H, Wilson R. 1986. The foundations of dynamic monopoly and the Coase conjecture. *J. Econ. Theory* 39:155–90

Haas E. 1953. The balance of power. *World Polit.* 5:442–77

Harsanyi J. 1967–1968. Games with incomplete information played by Bayesian players, parts 1–3. *Manage. Sci.* 14:159–82, 320–34, 486–502

Hicks J. 1932. *The Theory of Wages*. New York: Macmillan. 2nd ed.

Hinsley FH. 1963. *Power and the Pursuit of Peace*. New York: Cambridge Univ. Press

Hume D. 1898 (1752). *Essays: Moral, Political, and Literary*. New York: Longmans, Green

Huth P, Russett B. 1984. What makes deterrence work? *World Polit.* 36:496–526

Jervis R. 1976. *Perception and Misperception in International Politics*. Princeton, NJ: Princeton Univ. Press

Jervis R. 1979. Cooperation under the security dilemma. *World Polit.* 30:167–214

Kaplan M, Burns AL, Quandt R. 1960. Theoretical analysis of the balance of power. *Behav. Sci.* 5:240–52

Kennan J, Wilson R. 1993. Bargaining with private information. *J. Econ. Lit.* 31:45–104

Kim W. 1991. Alliance transitions and great power war. *Am. J. Polit. Sci.* 35:833–50

Kim W. 1992. Power transitions and great power war from Westphalia to Waterloo. *World Polit.* 45:153–72

Kleibor M. 1996. Understanding success and failure of international mediation. *J. Confl. Resolut.* 40:360–89

Knutsen T. 1997. *A History of International Relations Theory*. New York: Manchester Univ. Press. 2nd ed.

Kreps D. 1990. *Game Theory and Economic Modeling*. New York: Oxford Univ. Press

Kydd A. 1997. Game theory and the spiral model. *World Polit.* 49:371–400

Kydd A. 2000. Arms races and arms control. *Am. J. Polit. Sci.* 44:228–44

Labs EJ. 1992. Do weak states balance? *Sec. Stud.* 1:383–416

Lake DA, Powell R. 1999. International relations: a strategic-choice approach. In *Strategic Choice and International Relations*, ed. DA Lake, R Powell, pp. 3–38. Princeton, NJ: Princeton Univ. Press

Levy J. 1987. Declining power and the preventive motivation for war. *World Polit.* 40:87–107

Levy J. 1989. The causes of war. In *Behavior, Society, and Nuclear War*, ed. P Tetlock, J Husbands, R Jervis, P Stern, C Tilly, 1:209–333. New York: Oxford Univ. Press

Lynn-Jones S. 1995. Offense-defense theory and its critics. *Sec. Stud.* 4:660–94

Mansfield E. 1992. The concentration of capabilities and the onset of war. *J. Confl. Resolut.* 36:3–24

Mansfield E. 1994. *Power, Trade, and War*. Princeton, NJ: Princeton Univ. Press

Maoz Z. 1983. Resolve, capabilities, and the outcomes of interstate disputes, 1815–1976. *J. Confl. Resolut.* 27:195–229

Mattingly G. 1955. *Renaissance Diplomacy*. Boston: Houghton Mifflin

Mearsheimer J. 1990. Back to the future. *Int. Sec.* 15:5–56

Morgenthau H. 1967. *Politics Among Nations*. New York: Knopf. 4th ed.

Moul W. 1988. Balances of power and the escalation to war of serious international disputes among the European great powers, 1815–1939. *Am. J. Polit. Sci.* 35:241–75

Muthoo A. 1999. *Bargaining Theory with Applications*. New York: Cambridge Univ. Press

Myerson R. 1992. On the value of game theory in social science. *Ration. Soc.* 4:62–73

Niou E, Ordeshook PC. 1990. Stability in anarchic international systems. *Am. Polit. Sci. Rev.* 84:1207–34

Organski AFK. 1968. *World Politics*. New York: Knopf

Organski AFK, Kugler J. 1980. *The War Ledger*. Chicago: Univ. Chicago Press

Osborne MJ, Rubinstein A. 1990. *Bargaining and Markets*. San Diego: Academic

Powell R. 1994. The neorealist-neoliberal debate. *Int. Organ.* 48:313–44

Powell R. 1996a. Bargaining in the shadow of power. *Games Econ. Behav.* 15:255–89

Powell R. 1996b. Stability and the distribution of power. *World Polit.* 48:239–67

Powell R. 1999. *In the Shadow of Power*. Princeton, NJ: Princeton Univ. Press

Powell R. 2002. Game theory, international relations theory, and the Hobbesian stylization. In *The State of the Discipline*, ed. I Katznelson, H Milner. Washington, DC: APSA. In press

Roth A. 1979. *Axiomatic Models of Bargaining*. Berlin: Springer-Verlag

Rubinstein A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50(1):97–110

Sample G. 1997. Arms races and dispute escalation. *J. Peace Res.* 34:7–22

Schelling TC. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard Univ. Press

Schelling TC. 1966. *Arms and Influence*. New Haven, CT: Yale Univ. Press

Schultz KA. 1998. Do democratic institutions constrain or inform? *Int. Organ.* 53:233–66

Schultz KA. 2001. *Democracy and Coercive Diplomacy*. New York: Cambridge Univ. Press

Singer JD, Bremer S, Stukey J. 1972. Capabilities distribution, uncertainty and major power wars 1820–1965. In *Peace, War, and Numbers*, ed. B Russett, pp. 19–48. Beverly Hills, CA: Sage

Siverson R, Tennefoss M. 1984. Power, alliance, and the escalation of international conflict, 1815–1965. *Am. Polit. Sci. Rev.* 78:1057–69

Smith A. 1998. Fighting battles, winning wars. *J. Confl. Resolut.* 42:310–20

Ståhl I. 1972. *Bargaining Theory*. Stockholm: Stockholm School Econ.

Van Evera S. 1998. Offense, defense, and the causes of war. *Int. Sec.* 22:5–43

Wagner RH. 1986. The theory of games and the balance of power. *World Polit.* 38:546–76

Wagner RH. 1994. Peace, war, and the balance of power. *Am. Polit. Sci. Rev.* 88:593–607

Wagner RH. 2000. Bargaining and war. *Am. Polit. Sci. Rev.* 44:469–85

Walt SM. 1988. Testing theories of alliance formation. *Int. Organ.* 42:275–316

Walt SM. 1992. Alliance, threats, and U.S. grand strategy. *Sec. Stud.* 1:448–82

Walter BF. 1997. The critical barrier to civil war settlement. *Int. Organ.* 51:335–64

Waltz KN. 1959. *Man, the State, and War*. New York: Columbia Univ. Press

Waltz KN. 1979. *Theory of International Politics*. Reading, MA: Addison-Wesley

Weede E. 1976. Overwhelming preponderance as a pacifying condition among contiguous Asian dyads, 1950–1969. *J. Confl. Resolut.* 20:395–411

Werner S. 2000. Deterring intervention. *Am. J. Polit. Sci.* 44:720–32

Wittman D. 1979. How a war ends: a rational model approach. *J. Confl. Resolut.* 23:743–63

Wolfers A. 1962. *Discord and Collaboration*. Baltimore, MD: Johns Hopkins Univ. Press

Wright Q. 1965. *A Study of War*. Chicago: Univ. Chicago Press

Annual Review of Political Science
Volume 5, 2002

CONTENTS

ERRATA

An online log of corrections to *The Annual Review of Political Science* chapters (if any have yet been occasioned, 1997 to the present) may be found at http://polisci.annualreviews.org/