the continuation probability π and the lower the interest rate, the more heavily the future losses are weighted, and hence the greater is the prospect for cooperation. Second, while (*Unforgiving,Unforgiving*) can be a Nash equilibrium, it is by no means a unique equilibrium. A particularly well-known alternative strategy in the repeated prisoner's dilemma is *Tit-For-Tat* (*TFT*), whereby a player cooperates in the first round and then matches the other player's preceding move in each round thereafter. Using an argument similar to that noted earlier, it can be shown that (*TFT,TFT*) can be a Nash equilibrium with mutual cooperation in every round. Going to the opposite extreme, however, it can also be shown that both players' defecting in every round is also a Nash equilibrium, resulting in a decidedly Pareto inefficient outcome. Lastly, cooperation can emerge under standard assumptions only if the number of rounds is infinite or indefinite. In a game of known length, there is no incentive to cooperate in the end round and consequently, by backward induction, in any other round. Hence, for a finitely repeated game, mutual defection in all rounds is the predicted outcome.

4.4. Bibliographic Notes

This chapter is intended as a brief introduction to basic game theory, and as such there is much that is not covered. If we were to extend the introduction, we would want to include most particularly coverage of expected utility, mixed strategies, bargaining, incomplete information, and evolutionary games. Fortunately, there exist many fine texts that are rigorous and yet accessible. Among our favorites are Dixit and Nalebuff (1991), Gardner (2003), Dixit and Skeath (2004), and Binmore (2007), and, at a higher level, Kreps (1990) and Mas-Colell et al. (1995).

Because game theory's early impetus came largely from the Cold War, it is unsurprising that applications abound in conflict economics. For a single seminal contribution we would cite Schelling's (1960) wide-ranging analysis of rational behavior in situations that mix elements of conflict and common interest. For useful reviews of the application of game theory, see O'Neill (1994) on international relations and war, and Sandler and Arce (2007) on terrorism.

5

A Bargaining Model of Conflict

Economics involves the study of choices under conditions of scarcity. Whereas traditional economics assumes that choices are from among peaceful alternatives, conflict economics recognizes that some alternatives are violent or potentially violent. In this chapter we continue the move from traditional economics to conflict economics by sketching a simple economic model of conflict. The graphical model presented here is due originally to Hirshleifer (1995, pp. 172–175); more formal versions are available in Skaperdas (2006) and in Appendix B. Because the model is broadly consistent with what is known in political science as the bargaining (or rationalist) theory of war (Fearon 1995), we refer to it hereafter as Hirshleifer's bargaining model. Although by no means complete, the model provides a simple but effective framework for thinking systematically about the elements of conflict, some of the prominent explanations for war, and the possible effects of third-party intervention.

5.1. Elements of Conflict

Suppose that a disputed resource is to be divided between two players *A* and *B*. The players might be nation-states disputing territory, a government and a rebel group clashing over natural resources, or a government and a terrorist organization competing for control of a population. The players begin by diverting secure resources into arms, and then they divide the disputed resource either by fighting or by peaceful settlement. For simplicity, assume that each player chooses a level of arms that is the same whether fighting or settlement is anticipated. If the players fight, a portion of the disputed resource is destroyed, with the surviving portion divided between them based on their comparative arms and military technologies.

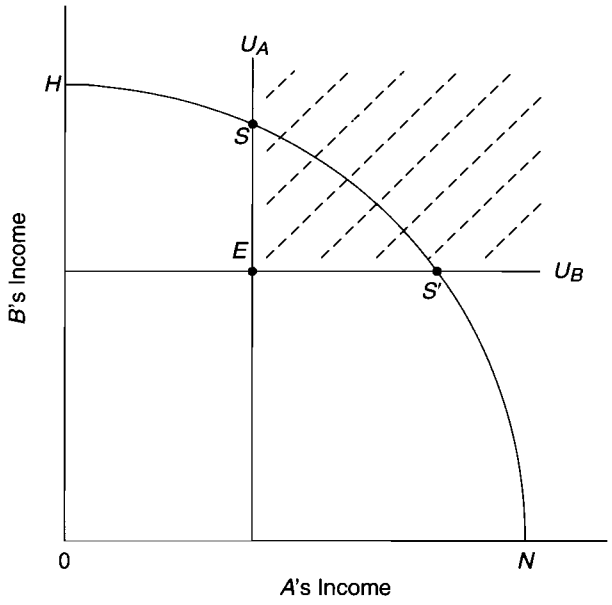

Figure 5.1. Bargaining model with peaceful settlement (adapted from Hirshleifer 1995, p. 172).

If they settle, the full amount of the disputed resource is divided by agreement, and the agreement is enforced by the threat embodied in the arms. Note here that the economic incentive to settle originates in the desire to avoid the resource destruction that accompanies fighting. On completion of either fighting or settlement, the players generate incomes from their final resources, which consist of their secure resources net of the diversion to arms plus their share of the disputed resource net of any destruction.

The choice between fighting and settlement rests on three elements: the expected income distributions based on fighting, the potential income distributions based on peaceful settlement, and the interpersonal preferences of the players. The attitudes of the players toward one another are reflected by their preferences over income distributions. These preferences may be benevolent, malevolent, or egoistic, depending on whether a player considers the other's income to be a good, a bad, or neither (called a neuter). Peaceful settlement occurs if there exists one or more distributions by settlement that would leave at least one player better off and neither player worse off than would be expected if they were to fight.

The elements of the model are illustrated in Figure 5.1 with a case of peaceful settlement. Assume that the players in the figure have complete information about armaments, conflict technology, and productive capabilities. This means that each player holds the same expectation about the outcome of any fight and resulting income distribution, shown by the conflict (or disagreement) point *E*. If the resource is divided peacefully, then alternative income distributions are feasible. These potential distributions are shown by the settlement-opportunity curve, which is labeled *HN*. Both players are assumed to be strict egoists, as indicated by their respective indifference curves U_A and U_B passing through point *E*. Player *A* cares only about her own income; thus her indifference curves are vertical, and she prefers all distributions to the right of *E*. For the same reason, player *B*'s indifference curves are horizontal, and he prefers all distributions above *E*. The highlighted area above and to the right of *E* is the region of mutual gain and includes distributions that are Pareto preferred to *E*, that is, distributions that make at least one player better off and neither player worse off relative to *E*. Because the settlement-opportunity curve intersects the region of mutual gain, the model predicts peaceful settlement over violence, resulting in a final outcome somewhere between points *S* and *S'* along the settlement-opportunity curve.

5.2. Sources of Violent Conflict

Inconsistent Expectations

Various elements of the situation depicted in Figure 5.1 combine to generate a peaceful settlement. To see how fighting might arise instead, assume now that the players' expectations are inconsistent. In particular, suppose in Figure 5.2 that player *A* expects the outcome of fighting to be income distribution E_A located toward the lower right, while player *B* expects it to be E_B toward the upper left. The region of mutual gain now lies entirely outside the settlement-opportunity curve *HN*. Because of the players' divergent expectations, the settlement-opportunity curve fails to intersect the region of mutual gain, and the predicted outcome is fighting.

Divergent expectations in this model imply that the players have incomplete information about various factors that determine the outcome of fighting, such as the other's capabilities, costs, strategies, or tactics. Figures 5.1 and 5.2 together suggest that an exchange of relevant

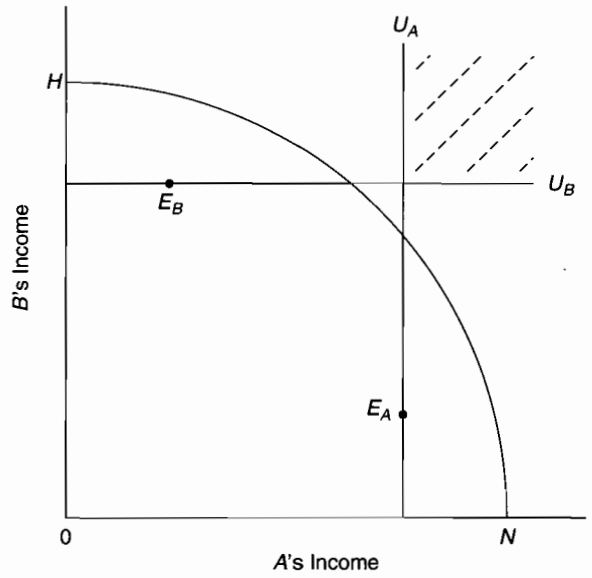

Figure 5.2. Inconsistent expectations with fighting (adapted from Hirshleifer 1995, p. 173).

private information between the players could generate a peaceful settlement that would be preferred to fighting. However, because expected outcomes determine which settlements are acceptable, the players have an incentive to withhold, exaggerate, or misrepresent certain information in order to attain a settlement more to their own advantage. This means that private information, even if accurate, tends to lack credibility when it is provided by the players. In this way, the exchange of information that might close divergent expectations is problematic (Fearon 1995).

Commitment Problems

A commitment problem exists when one player cannot trust the other because the latter has an incentive to renege after an agreement is reached. The essence of the problem is captured by the simplified game in Figure 5.3, which is part of a larger game wherein player *B* has proposed a specific settlement. Player *A* can either *Settle*, meaning agree to the peaceful settlement, or *Fight*. If *A* chooses *Settle*, then *B* can either *Abide* or *Renege*. Payoffs appear at the ends of the game-tree branches, with *A*'s payoff listed first. As shown, both players are better off by acting in accordance with the

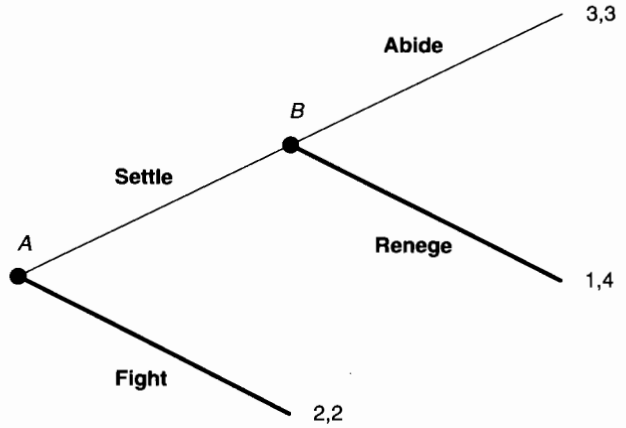

Figure 5.3. Two-player commitment problem.

peaceful settlement than by fighting, but given the settlement, *B* is better off by reneging than by abiding. Knowing that *B* is better off reneging, *A* mistrusts *B* and so chooses *Fight*. The predicted outcome is fighting with payoffs (2,2), despite its Pareto inefficiency when compared to the superior settlement with payoffs (3,3). To see how such commitment problems can lead to fighting, we look at three cases that are prominent in bargaining theory (Fearon 1995, Powell 2006).

Indivisibilities

Partial or complete indivisibility of a disputed issue is sometimes believed to be an important source of violence in conflicts over territory or political control (see Goddard 2006 and Toft 2006). Suppose two players dispute control of a sacred site that both players perceive to be completely indivisible. An example is shown in Figure 5.4, where each player is assumed to obtain income only when the site is completely held. If player *A* controls the site, the outcome is at point *N*; if *B* controls the site, the outcome is at point *H*. Thus the settlement opportunities frontier is represented entirely by the two points *H* and *N*. The expected outcome of fighting is shown by point *E*, which is determined by multiplying each player's probability of winning times her or his income from controlling that portion of the site not destroyed by fighting. Since the region of mutual gain includes neither settlement point *H* nor *N*, fighting is the predicted outcome.

Powell (2006) notes, however, that there are ways to peacefully distribute the disputed resource such that each player's expected payoff is

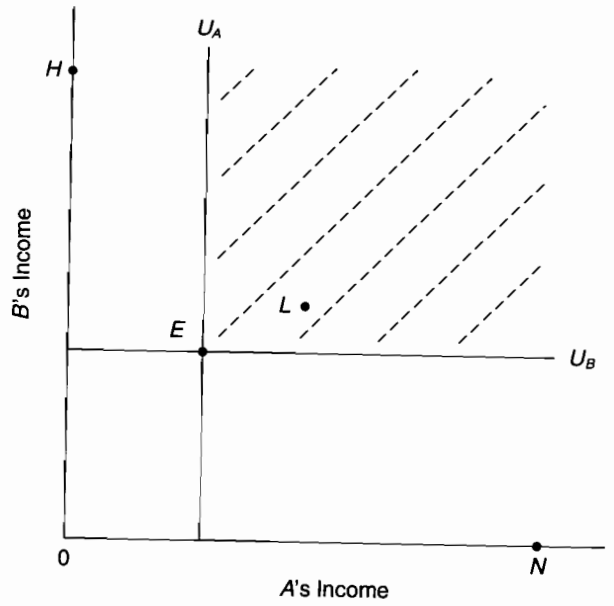

Figure 5.4. Commitment problem and indivisibilities.

greater under a settlement than under fighting. To see this, suppose each player has a 50 percent chance of winning control of the site through fighting, but each would suffer costs because of destruction. A lottery (e.g., a coin flip) that peacefully distributes the disputed item to each player with the same 50 percent probability would be preferred by both players because it avoids the costs of fighting. This is shown in Figure 5.4 by point *L*, which lies along a straight line (not shown) between *H* and *N*. Because costs of destruction are avoided under the lottery, point *L* lies within the region of mutual gain and is strictly preferred by each player to point *E*. Powell points out, however, that there is a commitment problem associated with the lottery. After the lottery, the disputed resource is distributed to the winning player, thus placing the outcome at either point *H* or *N*. But the loser has an incentive to renege on the lottery and initiate war, because the expected income from fighting at *E* is greater than the zero income assigned by the lottery.

Preemptive War

Consider next the case of preemptive war, which can arise from the existence of a first-strike advantage. Assume that the players have complete information and thus correctly anticipate the offensive advantage. In

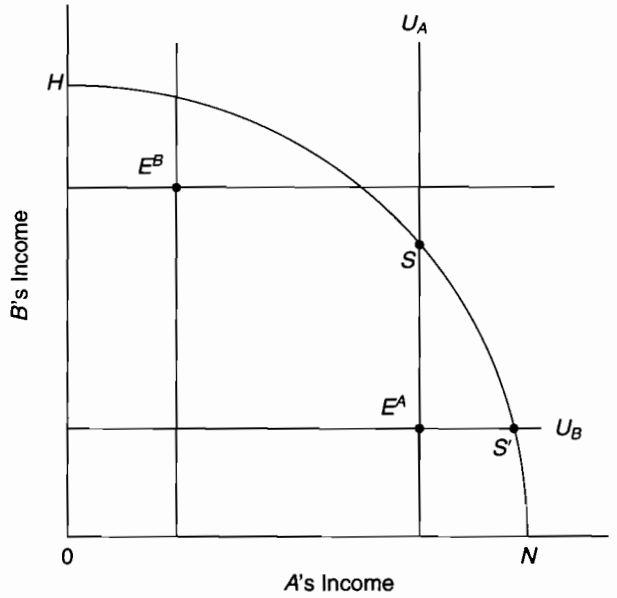

Figure 5.5. Commitment problem and preemptive war.

Figure 5.5, both players expect fighting to result in income distribution E^A if player *A* attacks first and in distribution E^B if player *B* attacks first. Suppose that the players bargain relative to E^A, the expected distribution if *A* attacks first. Any agreement by *B* to a settlement between *S* and *S'* is not credible, because *B* will have the incentive to renege, seize the first-strike advantage, and thereby shift the expected outcome to E^B. Knowing this, player *A* will be inclined to launch her own preemptive strike. By the same argument, any agreement by *A* relative to E^B is not credible, thereby leading *B* to strike. Because neither player can credibly commit to a mutually advantageous settlement, fighting is the predicted outcome.

Preventive War

To understand preventive war as a third case, the bargaining model must be extended to allow for dynamic considerations across time. Suppose the model consists of two periods. In period 1, the players choose levels of arms that result in conflict point *E* and settlement-opportunity curve *HN*. As seen in Figure 5.6, settlements between *S* and *S'* offer mutual gains relative to fighting. However, suppose it is known that a potential change in military technology exists that, if realized, would shift power toward player *B* in period 2. Assume that the shift in power can be prevented by

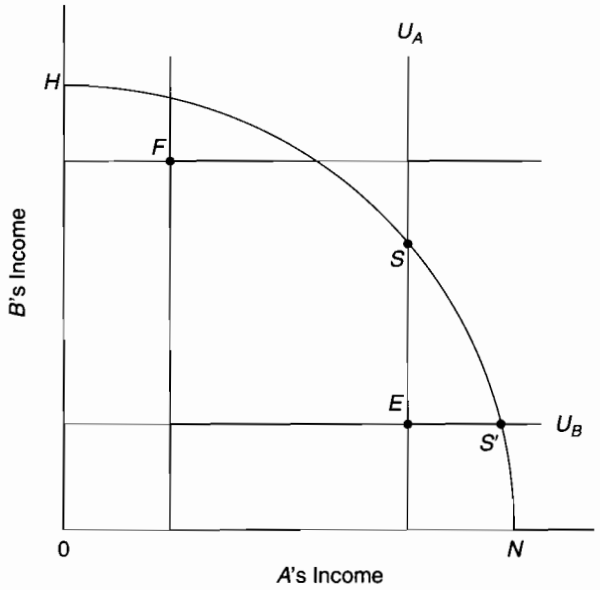

Figure 5.6. Commitment problem and preventive war.

fighting in period 1 but not by settlement. If the potential shift in power is sufficiently large, then a commitment problem arises, and fighting in period 1 is predicted.

To see the commitment problem most simply, assume in Figure 5.6 that arms levels (but not arms productivities) remain fixed, so that the change in technology would shift the conflict point northwest to point F in period 2 but would leave the settlement opportunity curve unchanged. In this case, an agreement by B in period 1 to a settlement between S and S' would not be credible. This is because such an agreement would leave B with a clear incentive in period 2 to threaten a fight and thereby secure a more advantageous settlement somewhere to the northeast of conflict point F. Knowing this, player A will weigh the prospective loss of income in period 2 against the cost of fighting in period 1. If the prospective shift in power is sufficiently large (as it is in Figure 5.6), player A will refuse settlement in period 1 and instead will lock in the distribution at point E by fighting.

Malevolent Preferences

As another source of violence, suppose that both players are highly malevolent. This means that they are willing to sacrifice large amounts of

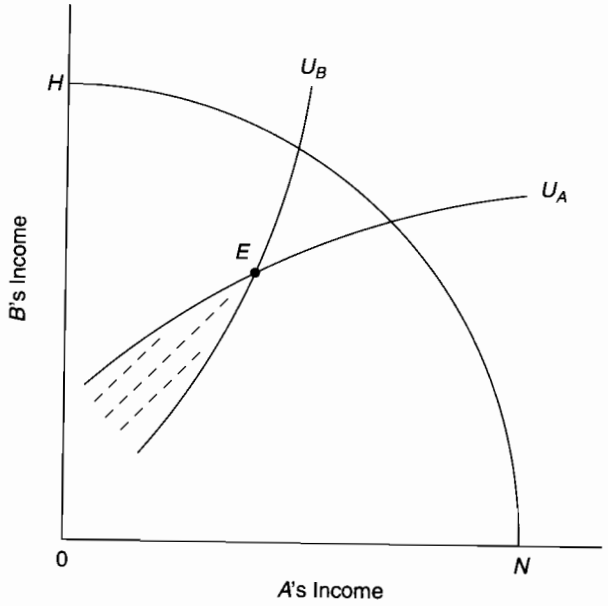

Figure 5.7. Malevolent preferences with fighting.

their own income in order to reduce the other's income. In Figure 5.7 the mutual malevolence of the players is shown by the positively sloped indifference curves passing through the conflict point E. Player A wants more income for herself and less for B; thus she prefers distributions down and to the right of E. Likewise, player B wants more income for himself and less for A; thus he prefers distributions up and to the left of E. In the figure, the region of mutual gain forms to the southwest of E and is nowhere intersected by the settlement-opportunity curve. This implies that the players would prefer to reduce each other's income further, but they can destroy resources only by fighting, not by settlement. Thus, to injure each other, the players can do no better than fight, and so fight they do.

Lesser levels of malevolence can leave open the possibility of peaceful settlement. This is shown in Figure 5.8, where the positive slopes of the indifference curves indicate mutual malevolence but nonetheless generate a region of mutual gain intersected by the settlement-opportunity curve. Notice, however, that the set of acceptable settlements SS' is diminished relative to the case with egoistic preferences in Figure 5.1. The general lesson is that malevolence reduces and possibly eliminates the range of settlements to which the players might agree.

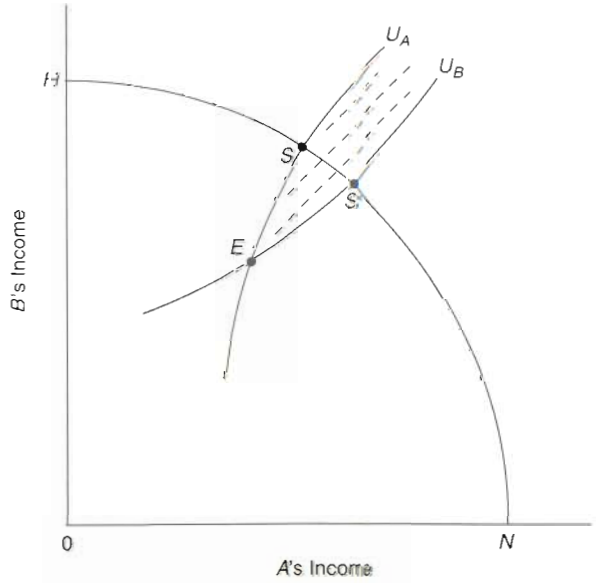

Figure 5.8. Malevolent preferences with peaceful settlement (adapted from Hirshleifer 1995, p. 173).

Political Bias

To this point we have assumed that players A and B are unitary actors who distribute a disputed resource by costly fighting or peaceful settlement. It is recognized in the social science literature, however, that a critical leader who substantially affects decisions to fight or settle can have different incentives relative to the group for whom the leader acts. Tarar (2006), for example, adapts Fearon's (1995) bargaining model of war to account for a critical leader's incentive to initiate war as a diversion from domestic problems. Also within a bargaining model, Jackson and Morelli (2007) consider a critical leader's "political bias," which encompasses anything that might cause the leader to have different incentives for war or peace relative to the group as a whole.

Here we relax the unitary actor assumption in Hirshleifer's model to illustrate how incentives facing a critical leader can be a distinct source of violence within the bargaining theory of war. For simplicity, suppose that only one of the players, say B, is subject to political bias. Assume also that each player chooses a level of arms that is unchanged whether fighting or settlement is anticipated and that any settlement achieved is enforced by those arms. Let b and $\hat{b}$ represent the proportion of B's income controlled

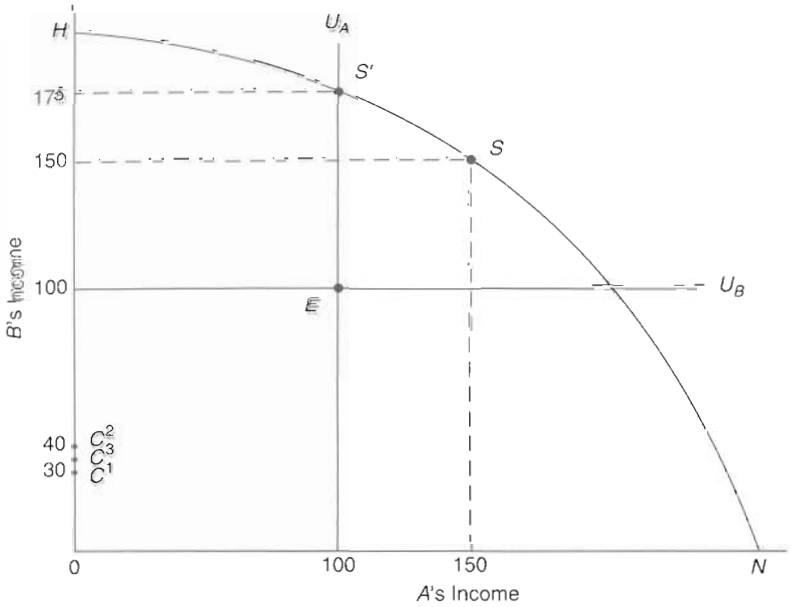

Figure 5.9. Possible fighting when player B is subject to political bias.

by the critical leader under peaceful settlement and fighting, respectively. For example, if $b = 0.2$ and $\hat{b} = 0.3$, then the critical leader in B controls 20 percent of B's income under peace but 30 percent of the income under fighting. Following Jackson and Morelli (2007, p. 1357), let $\hat{B} = \hat{b}/b$ represent the political bias of player B's critical leader. When $\hat{B} > 1$, the critical leader is biased in favor of fighting. For example, the critical leader might gain power or status with fighting that allows the leader to control a greater share of B's income than with settlement. Alternatively, a critical leader perpetrating atrocities during a war might expect retribution under settlement, causing the leader to control a smaller share of B's income under peace relative to fighting. Political bias does not necessarily favor fighting. When $\hat{B} < 1$, the pivotal leader controls a smaller share of income under fighting than under settlement. This might arise, for example, when peace confers popularity on the critical leader. When $\hat{B} = 1$, the leader's incentives between fighting and settlement match those of the broader group.

The essential trade-off for a critical leader with a political bias in favor of fighting is between controlling a larger share of less income by fighting versus a smaller share of more income by settling. To see how fighting might arise owing to political bias, begin by assuming in Figure 5.9 that

$b = 0.2$ and $\hat{b} = 0.3$ and that the critical leader controls the decision to fight or to settle. For players A and B as unitary actors, the outcome of fighting at E returns an expected income of 100 to each player. By avoiding the destructiveness of fighting, peaceful settlement at (say) point S promises each player an income of 150. Clearly, settlement at S is better for each player relative to fighting. In this particular case, the critical leader of B is indifferent between fighting and peaceful settlement at S: as shown in the figure, the leader will achieve the same income of 30 at point C^1 under fighting ($0.3 \times 100 = 30$) or peace ($0.2 \times 150 = 30$). Starting from point C^1, however, suppose that the proportion of B's income from fighting controlled by the critical leader rises from $\hat{b} = 0.3$ to $\hat{b} = 0.4$, other things constant. Now the critical leader for B will expect to gain more at C^2 from fighting ($0.4 \times 100 = 40$) than at C^1 from settlement ($0.2 \times 150 = 30$). Since the critical leader controls the decision to fight or to settle, fighting is the predicted outcome over settlement at S.

In a bargaining context, player A could attempt to appease the critical leader in B by offering a more generous settlement, such as point S' in Figure 5.9. As drawn, point S' reflects the maximum that A would be willing to offer the critical leader of B, because A's expected income from fighting would just equal her income from settlement at S'. Nevertheless, such a settlement offer would not be enough to prevent fighting. Under settlement at S', the critical leader of B would achieve income shown by point C^3 ($0.2 \times 175 = 35$), which is less than the leader's expected income of 40 at C^2 from fighting. Hence, the critical leader would still have an incentive to initiate a fight.

5.3. Third-Party Intervention

We have employed Hirshleifer's bargaining model to illuminate the elements of conflict and to identify several sources of violence. By extension, the model can also be used to suggest how third-party intervention might moderate violent conflict. Several of the various possibilities are shown in Figures 5.10 through 5.12. In each figure we assume for simplicity that the players' arms levels are unchanged by third-party intervention.

Consider first the possibility of economic intervention. In Figure 5.10, assume that a first-strike advantage exists such that the expected outcome is E^A if player A attacks and E^B if player B attacks. Consequently, the highlighted region of mutual gain lies outside of the settlement-opportunity curve, setting up the likelihood of preemptive war. Suppose, however, that a third party offers economic inducements to both players

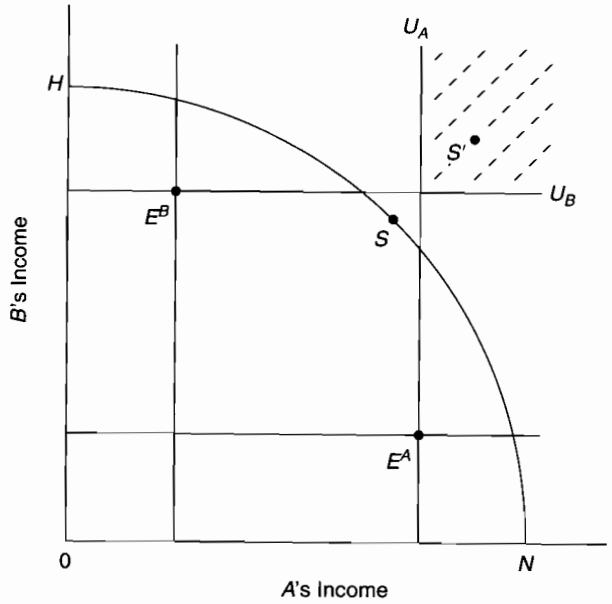

Figure 5.10. Third-party economic intervention.

contingent on peaceful settlement. If the inducements are sufficiently large and are tied to suitable settlements, then distributions within the region of mutual gain are feasible and settlement is expected. For example, conditioned on settlement at point S, the third party could offer subsidies to the players that would yield a final distribution at point S'.

As alternatives to economic intervention, a third party might provide diplomatic mediation (an example of peacemaking) or deploy military forces (an example of peace enforcement). In Figure 5.11, assume that A expects the outcome of fighting to be E_A while B expects it to be E_B. Since the highlighted region of mutual gain lies outside of the settlement-opportunity curve, fighting is the predicted outcome. Suppose, however, that a third party facilitates the credible exchange of private information between the players. If the expected outcomes of fighting are brought sufficiently close, then peaceful settlement becomes feasible. For example, if mediation shifts expectations to points $\tilde{E}_A$ and $\tilde{E}_B$, then in this case peaceful settlement between points S and S' along the settlement-opportunity curve is made possible. Alternatively, suppose that a third party undertakes military efforts to terminate violence between the players. In particular, assume that a third party deploys military forces against both A and B, thus eroding their expected returns from fighting. In this case, the expected

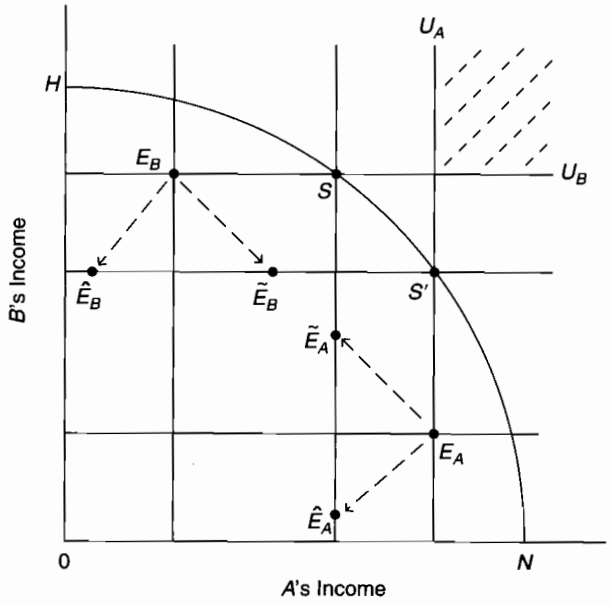

Figure 5.11. Third-party mediation or military intervention.

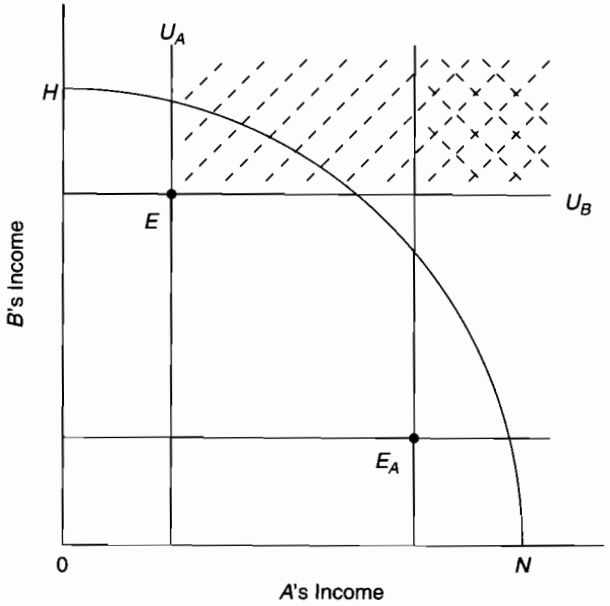

Figure 5.12. Third-party intervention favoring player *A*.

outcomes from fighting shift southwest to $\hat{E}_A$ and $\hat{E}_B$ in Figure 5.11, and peaceful settlement between S and S' is again made possible.

Figures 5.10 and 5.11 show how third-party intervention can ameliorate violence. If an intervention favors one side, however, it can leave the prospect for peaceful settlement unchanged or even worsened. For example, in Figure 5.12 assume that both players initially expect an outcome of fighting at point E, thus forming a region of mutual gain intersected by the settlement-opportunity curve. Now suppose that a third party intervenes privately to the advantage of player A, such that A expects the fighting outcome to be E_A while B continues to expect E. This means that the region of mutual gain, shown by the cross-hatched area, is reduced in size and now lies outside the settlement-opportunity curve. In this case, third-party intervention reduces the prospect for settlement.

5.4. Bibliographic Notes

The connection between conflict and bargaining was made by Schelling (1960), who famously wrote, "To study the strategy of conflict is to take the view that most conflict situations are essentially *bargaining* situations" (p. 5). The bargaining model sketched in this chapter is a modification of one first outlined by Hirshleifer (1985; 1987, ch. 10) and extended in Hirshleifer (1995) and Hirshleifer, Glazer, and Hirshleifer (2005). Early predecessors include Bush's (1972) model of anarchy and Wittman's (1979) model of war termination. Later models that are both more complete and more formal than Hirshleifer's model include Skaperdas (2006) and Garfinkel and Skaperdas (2007).

Whereas we assume that the contested prize is a resource, Hirshleifer (1995) leaves the nature of the prize unspecified. How the contested prize is specified constitutes a major distinction among formal models of conflict. In particular, the object up for grabs can be defined alternatively as a resource (Hirshleifer 1991, Grossman and Kim 1995, Anderton, Anderton, and Carter 1999), a production good (Rider 1999, Hausken 2004), a consumption good (Bush 1972, Bush and Mayer 1974), or an export or import (Anderton and Carter 2008b).

Various sources of violent conflict are explored individually in this chapter, but wars can have multiple causes or antecedents (Vasquez 2000, Levy 2008). Despite complexities, several sources of violence could be modeled together in the Hirshleifer framework. Other potential sources of violence not explored in this chapter include players' concerns for reputation (Schelling 1966, Walter 2006a, Crescenzi, Kathman, and Long

2007), risk-seeking preferences and increasing returns to production (Skaperdas 2006), and an intertemporal incentive to weaken or eliminate a rival (Garfinkel and Skaperdas 2000a, 2007).

As noted in the chapter's opening, Hirshleifer's model is broadly consistent with the bargaining theory of war in the political science literature. For literature reviews, see Powell (2002) and Reiter (2003). Particularly instructive are the costly lottery models of war found in Fearon (1995) and Powell (2006). For a game-theoretic discussion of commitment problems, see chapter 10 of Dixit and Skeath (2004). For formal conflict-settlement protocols, see Isard and Smith (1982), Raiffa (1982), Brams and Taylor (1996), and Garfinkel and Skaperdas (2007, pp. 667–682). Models of third-party intervention include Amegashie and Kutsoati (2007) and Chang, Potter, and Sanders (2007).

6

Conflict between States*

For millennia, philosophers and sages have pondered the origins and horrors of war. Despite this long history of inquiry, it is only in the last century that scholars from political science, economics, and other disciplines have attempted to use the quantitative methods of social science to study the causes and effects of war. Building on the early work of Lewis Richardson, Pitirim Sorokin, and Quincy Wright, the social scientific study of war was well established by the mid-1960s around a community of scholars associated with the Correlates of War Project, the Peace Science Society (originally, Peace Research Society), the *Journal of Conflict Resolution*, and the like. Since then, a wealth of social scientific studies of war has appeared in journals and books across the various disciplines (Singer 2000, Anderton and Carter 2007). In this chapter we focus on armed conflict between states, before we turn to civil war in Chapter 7.

6.1. The Conflict Cycle

Conflicts typically pass through phases, as shown by Lund's (1996) life-cycle diagram in Figure 6.1, which plots the level of conflict between parties across time. The conflict in question may be interstate as covered in this chapter, intrastate as in Chapter 7, or extra-state as in Chapter 8. The bell-shaped curve represents the course of a typical conflict as hostility rises and falls over time. The vertical axis marks levels of conflict beginning with durable peace and rising successively to stable peace, unstable peace, crisis,

* The introductory paragraph, sections 6.1 and 6.2, and parts of section 6.4 of this chapter are adapted from Charles H. Anderton and John R. Carter, "A Survey of Peace Economics," published in *Handbook of Defense Economics*, volume 2, edited by Todd Sandler and Keith Hartley, pp. 1211–1258, Copyright © Elsevier 2007. We gratefully acknowledge Elsevier's permission to republish material from the article.