

War and the State

The Theory of International Politics

R. Harrison Wagner

THE UNIVERSITY OF MICHIGAN PRESS

Ann Arbor

For Rosi

Copyright © by the University of Michigan 2007

All rights reserved

Published in the United States of America by

The University of Michigan Press

Manufactured in the United States of America

⊗ Printed on acid-free paper

2010 2009

4 3 2

No part of this publication may be reproduced,
stored in a retrieval system, or transmitted in any form
or by any means, electronic, mechanical, or otherwise,
without the written permission of the publisher.

A CIP catalog record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data

Wagner, R. Harrison (Robert Harrison)

War and the state : the theory of international politics /

R. Harrison Wagner.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-472-09981-8 (cloth : alk. paper)

ISBN-10: 0-472-09981-7 (cloth : alk. paper)

ISBN-13: 978-0-472-06981-1 (pbk. : alk. paper)

ISBN-10: 0-472-06981-0 (pbk. : alk. paper)

1. International relations—Philosophy. 2. War. 3. State, The.
I. Title.

JZ1316.W34 2007

327.101—dc22

2006028159

ISBN13 978-0-472-02590-9 (electronic)

Violence, Organization, and War

What is a war? That is a question that is harder to answer than many people assume. At a minimum it is a contest in which organized groups compete in killing and wounding each other or destroying things they value. But why would any group expect to benefit from harming members of another group?

Often such contests take the form of an attempt by one group to destroy the military forces of another, and many people tacitly, but incorrectly, assume that that is the point of all military contests. But what would be the benefit of doing that, even if one were successful? A common answer is that it would prevent the other group from using its military forces against one's own. But that only raises the question of why the other group would profit from doing it: if it is common knowledge that no group would benefit from disarming another group, then no one could expect to profit from destroying another group's means of disarming one's own.¹

The effect of one group's disarming another is just that the victorious group can then kill or wound the members of the other group, or destroy things of value to them, without forceful opposition. Moreover, contrary to what is commonly assumed, in many violent contests neither side has any prospect of destroying the military forces of the other side, and even if they did, neither seems to try seriously to do it. Thus our primary task must be to explain why one group would expect to benefit just from killing or wounding the members of another group or destroying their property.

We saw at the beginning of chapter 1 that John Mearsheimer (2001) claimed that international politics was "tragic." What he meant was that war is inefficient, though admittedly the word *inefficient* seems hardly adequate as a way of describing war. I use the word *inefficient* here in the way that it is used in economics, where it describes a consequence of the choices of two or more people that leaves them worse off than they would have

1. Note that that would appear to describe a world of Hobbesian commonwealths, in which it would be common knowledge that the function of sovereigns, to use Hobbes's words, was simply to provide security to their citizens from the "injuries of one another" and "from the invasion of foreigners."

been had they chosen differently.² To say that war is inefficient means that there is an alternative to war that would be better for the participants than the prewar expected value of fighting.³

But why would a war ever occur if all the parties to it would be better off avoiding it? One possible answer is that strong emotions or cognitive errors of some sort have prevented them from seeing their true interests, and this seems to be what is implicitly assumed by much of the literature on conflict management. However, game theory provides compelling examples of situations in which rational individuals choose outcomes they would all have been better off not choosing. One is the coordination game discussed in the previous chapter. Two others, the Prisoner's Dilemma and the Stag Hunt, have been especially popular as ways of understanding why much of what happens in international politics seems to be contrary to everyone's interests, as we saw in chapter 1. However, all these games are far more plausible as explanations of a failure to cooperate to achieve common interests than as explanations of war. And if people are unable to cooperate, there can be no armies and therefore no wars.

Kenneth Waltz wrote, "The threat of force internationally is comparable to the role of the strike in labor and management bargaining" (1979, 114). This suggests another explanation of inefficient choices that is a much more promising way of understanding wars than the Prisoner's Dilemma or the Stag Hunt but one that Waltz never developed.

In a strike an organization of employees seeks to increase employees' gains from a contract by preventing the owners from profiting from the firm or industry until they agree to terms that are more favorable for the employees. However, in doing this the employees also harm themselves. Thus whatever the terms of the agreement that ends a strike, the strike itself is costly for both sides, and both would have been better off accepting the agreement before the strike rather than after it. This is often true of wars as well.

Strikes are commonly regarded as examples of bargaining, and therefore one might hope that an understanding of bargaining would help one understand why strikes occur. If Waltz is right, then an understanding of bargaining might also contribute to an understanding of wars.

Strikes, of course, may be accompanied by violence. However, there are two attributes of strikes that distinguish them from most violent

2. More precisely, an outcome is inefficient if there is another feasible outcome that at least one person would prefer and that would leave no other persons worse off.

3. The classic discussion of the tragic nature of interstate conflict is Butterfield 1951. But Butterfield's main example is the struggle between East and West over the future of Germany in the aftermath of World War II, which did not lead to war. It is instructive to read Marc Trachtenberg's (1999) discussion of this issue in the context of Butterfield's analysis. See also Wagner 1980.

conflicts, and if we are to understand violent conflicts we must bear these differences in mind. One is that even if force is used by labor or management, a wage contract leaves them both better off than they would be without one. In the context of war, however, force is used in an attempt to make one's adversary worse off. Second, wage contracts are typically enforceable, whereas the parties to an agreement made in the context of the use of force must usually be concerned about whether the agreement will in fact be carried out.

These two differences are more closely related to each other than they may at first appear to be. Whether some contract makes one party better off or not depends on the baseline used for comparison. In adverse economic circumstances management may succeed in getting labor to agree to a wage agreement that leaves workers worse off than they were before, and the workers may therefore feel no different from someone who surrenders something valuable to someone else at gunpoint. Distinguishing between the two cases requires a distinction between what one possesses and what one owns. But this distinction rests on a definition of property rights, which may be contested and in any case has to be enforced.

We should not exaggerate the enforceability of contracts even when the definition and enforcement of property rights can be taken for granted. Wage bargains often include complex stipulations concerning working conditions and management prerogatives that can be evaded without provoking external sanctions. However, bargaining theory takes the enforcement of agreements for granted, and if we are to understand violent conflicts we must consider what the consequences might be if agreements can be violated. First, however, we must think about how to understand bargaining.

Bargaining

A strike is, at least in part, a dispute between labor and management about how to divide up the revenues of a firm. Thus it can be helpful in thinking about strikes to think first about a simpler situation in which two people are offered a sum of money if they can agree on how to divide it between them, but if they cannot agree they get nothing. We can ask two questions about such a situation: (1) What division will the two bargainers agree to? and (2) How long will it take them to agree?

These simple questions lead to two surprisingly difficult puzzles, which are the subject of a very large literature. The first is the result of the fact that rationality and self-interest alone are not enough to answer the question of what division the bargainers will accept. The second puzzle concerns the relation between the two questions just stated: if it is clear

what division the bargainers should accept, then they should accept it immediately. But if that were true then bargaining as it is commonly understood (including strikes) would never occur.

Why Haggle?

Since it seems plausible that rational bargainers will not throw the money away and therefore will certainly agree to something, the question of what they will agree to may seem unimportant. However, the question of how long it will take them to agree is very important, since even temporary disagreement can be extremely costly. A plausible answer is that agreement on a division of the money requires some information that the two bargainers may not initially possess, and if they do not have it the bargaining process provides a way of getting it. Thus the costs associated with delay in reaching agreement are the price that must be paid for the information required to reach it. However, an understanding of what that information might be requires a solution to the first puzzle, which is therefore more important than it first appears to be. Let us see how it arises.

Since by assumption each bargainer prefers more money to less but neither will get any unless the other agrees, each bargainer's decision about what division to accept depends on her expectation of what the other will agree to. This is the sort of problem that game theory was invented to solve, and the answer to it was supplied by John Nash and is therefore called the Nash equilibrium. But all that the Nash equilibrium requires of rational bargainers is that their expectations be consistent, in the sense that, given some expectation as to what each will do, neither should have an incentive to deviate from it. And if each bargainer is free to demand any amount as a condition for his or her agreement then every possible division of the money satisfies this requirement. Thus game theory seemed at first merely to justify the common belief that any division of the money would be consistent with rational behavior, and therefore nothing could be said about what rational bargainers would agree to.⁴

Two types of solution to this puzzle have been offered.⁵ One is that, since every division would be a Nash equilibrium, the problem is, as in the coordination game discussed in the previous chapter, to coordinate the

4. For an example of a Nash equilibrium, see the discussion of the coordination game in chapter 2. It is important to distinguish the Nash equilibrium, which is a necessary condition for rational behavior when decisions are interdependent, from the Nash bargaining solution, which was Nash's own attempt to answer the question of what bargainers should agree to. Unfortunately the Nash bargaining solution requires special axioms whose relation to individual rationality is by no means clear. Thus not even Nash thought that the Nash bargaining solution was a definitive answer to the bargaining problem.

5. For a more extended discussion and references to the literature, see Kreps 1990, 551–71.

bargainers' expectations on one division rather than another. This implies that the bargaining problem is at its core just another example of a coordination problem, though one that is complicated by the fact that the bargainers have conflicting preferences as to which division they coordinate on. And therefore salient divisions, conventions, or prevailing conceptions of fairness may lead the bargainers to focus their expectations on one particular outcome, and the costs associated with a failure to coordinate may deter them from deviating from what is expected. If two people are bargaining over the division of a sum of money, all these factors may lead them to coordinate their expectations around an agreement that divides the money equally.⁶

This answer to the puzzle implies that there are two possible explanations for a failure to reach immediate agreement. One is that the bargainers have failed to coordinate on a particular division of the money. Unfortunately it is not clear what they should do in this case. The other possible explanation is that they agree on how the gains should be divided but they disagree on how to measure them. This explanation is consistent with the suggestion that delay in reaching agreement is the result of a lack of relevant information.

The other solution to the puzzle posed by the existence of multiple Nash equilibria has two parts. One is to note that the definition of rational behavior given by the Nash equilibrium is incomplete, since in many situations it is consistent with behavior that is patently not rational. For example, suppose one of our bargainers thought he might gain if he threatened to detonate a bomb killing both bargainers if the other did not agree to his terms. If the other bargainer rejected his demand then the one who had made the threat would not want to carry it out, since the only consequence of carrying it out would be that he was killed along with the other bargainer. In other words, once the other bargainer refused his demand, carrying out the threat would no longer be part of a Nash equilibrium. A tighter definition of rational behavior would rule out equilibria that contained such incredible threats, and such a definition provides one part of a possible solution to the puzzle of too many Nash equilibria in bargaining situations.

The other part of the solution is to require that the process by which offers and counteroffers are made be modeled explicitly and any agreement then be the result of a (suitably refined) equilibrium combination of strategies in such a negotiation game. In modeling the bargaining process it is plausible that negotiators will prefer agreements that come sooner to

6. Note, however, that since the bargainers have conflicting preferences as to what division they should coordinate on, any convention that determined that might be the source of significant conflict. (Compare the example of the choice of language mentioned in the preceding chapter.)

agreements that come later, and even if a demand by one bargainer is accepted by the other an exchange of offers will require a finite amount of time. Ariel Rubinstein (1982) was the first to show that with these assumptions there is a unique combination of strategies that satisfy the requirement of sequential rationality just mentioned.⁷

Because both bargainers prefer present agreements to future ones, Rubinstein's model implies that they should reach agreement immediately. Thus explanations of delay focus on the possibility that some of the information required by the bargaining solution is missing. In Rubinstein's model an obvious candidate for this role is the discount rates of the two bargainers. Since each bargainer presumably knows her own discount rate, the problem must be that this information is not common knowledge. And since each bargainer has an incentive to misrepresent it, this problem cannot be overcome simply by having each bargainer reveal it to the other. Thus the only way each can acquire information about the other's discount rate is through observing what offers each makes and rejects in the course of the bargaining process, which provides a formal justification for the idea that the bargaining process allows for the revelation of information.⁸

If we are to use this reasoning as a way of explaining strikes, we would have to distinguish between prestrike exchanges of offers and exchanges of offers once the strike has begun. Any exchange of offers prior to a strike takes place while the firm is operating and therefore while labor and management are benefiting from some existing division of its revenues. If one is satisfied with that division but the other is not, then the fact that the satisfied party discounts future benefits provides it with no motivation to agree to any change in the status quo. It is rather the expected outcome of bargaining in the context of a strike that might motivate the satisfied party to agree to make some concession.

In bargaining theory the set of possible agreements is commonly called the bargaining frontier, and the outcome that would occur in the absence of agreement is called the disagreement outcome. Thus in prestrike negotiations labor or management threatens to revert temporarily to the disagreement outcome in order to renegotiate the terms of the wage bargain. However, if they share enough information about the consequences of doing so then this will not be necessary.

In prewar crisis bargaining, the disagreement outcome is war. In the

7. See the discussion of Rubinstein's argument in the next chapter.

8. Rubinstein's answer to the bargaining problem depends not just on a refinement of Nash's definition of rational behavior but also on Rubinstein's specific assumptions about how bargainers are expected to negotiate with each other. While these assumptions are not implausible, they are not the only plausible assumptions one might make. For a survey of the literature on this subject, see Kennan and Wilson 1993.

next chapter we will see that this analysis of bargaining in the context of strikes has important implications for understanding the occurrence of war. But for war to occur there must be organizations that expect to benefit from the use of force. Thus we must first consider what bargaining theory tells us about the benefits from organizing and from using force.

Bargaining and Organization

With complete information the outcome of Rubinstein's bargaining game is not an equal division of the money. Rather, the bargainer who gets to make the first offer can take advantage of the fact that her adversary discounts future benefits and demand a larger share. Thus Rubinstein's analysis seems appropriate for situations in which there is no commonly accepted norm concerning how the money is to be divided, but each is simply out to get as much as she can from the other. Because of Rubinstein's assumptions about how the bargaining process proceeds, each bargainer is able, in effect, to deliver a little ultimatum to the other: accept my demand now or pay the price of waiting until I consider yours later. The longer the other bargainer has to wait, the bigger the premium the one making the first offer can extract.⁹

In the limit one bargainer might be able to confront the other with a choice between accepting her demand or getting nothing at all. A bargainer able to deliver such an ultimatum (or take-it-or-leave-it offer) could successfully demand all (or nearly all) the money.¹⁰

Take-it-or-leave-it demands are usually not credible because they imply that the person making the demand would prefer no agreement at all to an agreement on any other terms, which is normally not true. However, such demands may be credible if one individual has many alternative bargaining partners. Then if one potential partner rejects a demand there are others to replace him. That is how an organization increases the bargaining power of workers: it prevents management from making many take-it-or-leave-it offers to individual workers.

But if a bargain struck between management and a labor union benefits all workers, then individual workers may have no incentive to contribute to the support of the union, since if others contribute a person who does not will benefit anyway and if others do not contribute one person's contribution would be ineffective. Thus workers may face what is

9. The Rubinstein bargaining model is discussed further in the next chapter.

10. If what is at stake is not the division of a commonly known sum of money but, for example, the price for which something will be sold, a seller able to make a take-it-or-leave-it offer may demand more than the other is willing to pay and there will be no bargaining process through which he can learn that he was mistaken. In that case it is possible that no mutually beneficial agreement will be reached.

known as a collective action problem in capturing the potential gains from bargaining with management. If so, their bargaining power will be less than it would otherwise be.¹¹

Bargaining and the Use of Force

If the delay between offers in the bargaining process modeled by Rubinstein is small enough, then the Rubinstein solution will deviate only slightly from the equal division that a norm of fairness might prescribe. However, a requirement that the bargainers receive equal benefits from agreement can have surprising implications. Suppose, for example, that the sum to be divided is one hundred dollars and that some benefactor has offered to pay a bonus of fifty dollars to one of the bargainers if agreement is reached. Then if the money is divided equally one bargainer will receive fifty dollars and the other one hundred dollars, and the gains from agreement will therefore be unequal. If the bargainers are to benefit equally, therefore, the bargainer whose gains will be supplemented must get only twenty-five dollars of the money to be divided and the other must get seventy-five dollars.

Moreover, if, instead of supplementing one bargainer's gains, someone is expected to take some action that will cost him fifty dollars if no agreement is reached, the effect is exactly the same: that person will gain from the agreement both the share of the money he receives and the fifty dollars he would have lost if no agreement had occurred. Thus if the two bargainers' gains are to be equal, he must receive only twenty-five dollars of the money to be divided. This example illustrates the fact that there are two normative issues raised by bargaining: how the gains from agreement should be divided and what disagreement outcome should be taken as the baseline from which the gains are measured. When we consider that the person who is expected to deprive one of the bargainers of fifty dollars in the event of no agreement may be the other bargainer, it also helps us understand one of the uses of force.

While there may be people who derive utility directly from harming others, most harm is done because it is a way of achieving some other benefit. For example, if someone occupies a piece of land that I want, I may kill him in order to take it. Instead of killing him, however, I could allow him to continue to work the land and threaten to kill or harm him if he refused to give me any food he produces above what is required to keep him alive.

11. Note that the incentives that give rise to the collective action problem resemble those in the Prisoner's Dilemma game. The seminal work on the collective action problem is Olson 1965. See also Hardin 1982.

If all incentives to use force were like the first, then violent conflicts would all be like conflicts between animal predators and their prey, and its point would always be simply to separate people from things of value that they control. Often, however, the point of violence for human predators is, as in the second example, to influence the behavior of the victim, which would be true even if all I wanted was to persuade him to give up what he had. But then it is unclear why harm would actually be done to him or his property. The literature on bargaining provides a possible answer to that question.

Seen in that context, force is a way by which individuals can manipulate the disagreement outcome in a bargaining situation in order to gain something at the expense of others. However, as in any bargaining situation, both would have an interest in avoiding its actual use. I might, for example, threaten to beat the man unless he agreed to share his harvests with me. But beating him prevents him from working. We will both gain, therefore, if I stop beating him and he begins working. Thus a situation in which the man is continually beaten and does no work is the disagreement outcome in a bargaining situation in which he and I negotiate the terms on which he will work for me, and beating him may be a way of revealing information about the relative gains from agreement.¹²

Given some expectation as to how the gains from agreement are to be divided, I have an interest in maximizing the other person's gains by minimizing the expected value to him of disagreement. However, some threatened consequences of disagreement may not be credible. For example, I might threaten to kill the person I want to work for me if he refuses to comply, but if he refused and I killed him I could never benefit from his work. Thus threatening to beat him is more credible than threatening to kill him.

However, if there are many alternative workers, anyone that I kill might be replaced by another. This is another example of how the existence of many alternative bargaining partners can make take-it-or-leave-it demands more credible and therefore strengthen a person's bargaining power.

Like a firm dealing with many individual consumers, an organized group can make take-it-or-leave-it demands of many individuals, and thus the potential gains from the forcible redistribution of possessions or the forcible exploitation of the labor of others provides a motivation for the organized use of force. Indeed, such organizations are sometimes spoken of as though they were firms selling a product for profit and sharing the proceeds among their members. The "product" of such an organization is protection, and what it "sells" is protection against itself.

12. For a discussion of bargaining between master and slave, see Berlin 1998. See also Morgan 1999.

Like business firms, such organizations create the potential for three types of conflict: conflicts with their “customers,” conflicts with competing “firms,” and conflicts within the organization over the division of its revenues. Bargaining with the use of force can occur in all these contexts.

Members of an organization of economic predators have conflicting interests in dividing the gains from predation, and since the gains from any redistribution can be shared among the members of any group that objects to the current distribution, the leader of every organization of economic predators has to be concerned about the emergence of a competing coordinator among his followers. A combination of punishment of individual dissidents with attempts to inhibit free communication among them can make opposition seem risky and preserve the lion’s share of the gains for the leader.

“Customers” could increase their bargaining power if they were organized, since then the exploiters could not make take-it-or-leave-it demands of individuals. However, if any agreements they reach would benefit exploited individuals whether they resisted or not, then, unlike competitors to the leader from within the organization but like workers without a union, they might have to overcome a collective action problem if they are to organize.

Thus the incentives faced by the “customers” of such an organization may resemble the Prisoner’s Dilemma, while the incentives faced by potential members of a protection organization resemble the Stag Hunt—they have an incentive to coordinate their expectations on a cooperative outcome. Indeed, both the name *Stag Hunt* and the supporting story from Rousseau are singularly appropriate, since the skills humans developed for hunting and herding large animals may well have facilitated the hunting and capture of other humans. Thus the differential incentives faced by economic predators and their victims would seem to favor predation.

However, there are three important limits to the gains from economic predation. One is that economic predators do not produce anything and are therefore dependent on their prey to produce the goods that will support them both. Of course, the potential prey may not produce enough to make predation more attractive than doing something productive. With increasing productivity, however, predation becomes profitable, but the predator becomes dependent on the prey for his own well-being.

The second potential limit on predation is that if one group of predators can organize then so can another, and therefore the gains from predation invite competition from other predators. This provides another, secondary, way that force becomes a means to an end, since one organized group of predators can use force to eliminate another and gain exclusive access to its “customers.” It can do that by destroying the competing organization’s instruments of coercion, disrupting it so that it can no longer

function as an organized group, or threatening to punish its members if they do not agree to go out of business. But as we saw in the previous chapter, one of the insights of the *raison d'état* literature is that recurring conflicts among predators can increase the bargaining power of their prey and thus over time reduce the gains from predation.

This, then, is one possible answer to the question we started with, of why one group would gain from forcibly disarming another.¹³ However, forceful contests between competing predators are risky and costly, and thus the leaders of two competing organizations may both prefer to reach an agreement rather than fight a contest in which each tries to eliminate the other. Thus bargaining is also relevant to understanding violent contests between predators.¹⁴

There are two types of agreement they might reach to avoid competing with each other: they can merge and share the revenues from extortion, or they can divide the "market" between them. In principle there are many ways they might agree to divide the "market," but dividing their "customers" geographically is obviously the most efficient way of organizing coercion, and the most efficient division of territory is into contiguous blocks.

Thus an understanding of bargaining over these two types of agreement seems directly relevant to understanding the development of the European state system. A set of agreements among economic predators to divide the world among them might be called a world of independent states, and one of the central questions raised by the literature on the European state system is why such a set of agreements could not be permanent or, if not permanent, at least renegotiated without the actual use of force.

The third important limit to economic predation is that it is possible for its actual or potential victims to overcome the collective action problem that they face and engage in forceful bargaining with the predators. The Prisoner's Dilemma and Stag Hunt games are helpful metaphors in thinking about how this might be possible. However, we must avoid being seduced into thinking that they might be models of any actual situation, which would involve many people with many choices and in which uncertainty about what people's preferences actually are would likely play an important role.

The reason for thinking that the victims of economic predation might

13. Note the irony of this discussion in the context of structural Realism: Jervis (1978) argued that substituting the Stag Hunt for the Prisoner's Dilemma provided a possible solution to the security dilemma and therefore might prevent war. But because the incentives of economic predators resemble those of the Stag Hunt rather than the Prisoner's Dilemma, predation is profitable and violent conflicts among predators can be worth their costs.

14. The implications of this point will be explored in the next chapter.

face a situation analogous to the Prisoner's Dilemma is that resistance is likely to be both dangerous and costly, anything that weakens the predator would benefit all his victims whether they cooperate in opposing him or not, and individuals acting alone could not expect to accomplish anything. People who cooperate in replacing the predator, however, could expect to benefit from doing so in a way they otherwise would not. That is perhaps why economic predators often have more to fear from within their own ranks than from the population they exploit.

Of course, competing predators can come from within the exploited population as well as from within the established predator's own ranks, and therefore one way of organizing resistance to an established ruler is to organize competing exploiters from among the exploited. If such a group is successful, however, the result may be either long-lasting violent competition among predators or the creation of an even more effective predator organization to replace the original one. In the first case, the gains from predation will be limited while the violence lasts, but the exploited will suffer nonetheless, and therefore life will be "solitary, poor, nasty, brutish, and short." In the second case the gains from predation will be increased rather than diminished. But neither case is fundamentally different from the examples already discussed.

To see why people who merely wanted to resist predation might not actually face a collective action problem, let us look again at the sources of the "dilemma" in the Prisoner's Dilemma. They are twofold: the preferences of the actors and the constraints under which they must choose.

Consider first why the preferences of the actors might be different. One possibility is that if enough other people cooperate, then the actions of individuals might not be entirely without effect. If so, then individuals might prefer to cooperate in resistance, if they were sufficiently confident that others would cooperate as well. Then their preferences would resemble the preferences in the Stag Hunt game rather than the Prisoner's Dilemma.

Now consider the constraints on the choices of the prisoners in the Prisoner's Dilemma: they must independently choose between only two alternatives and do so only once. If their choices are repeated, or if they have more choices and do not choose independently, then even if their preferences remain the same they may no longer have a single dominant choice, and coordination on a mutually beneficial outcome may be possible.

For example, just as people who expect to do business repeatedly with each other have an incentive in any particular transaction not to cheat, so people who live in the same village may have an incentive to support a villager who is being treated in a way they consider unfair, if they want support from others when they are treated unfairly. Moreover, everyone

engaged in long-term relations with others has at his disposal ways of rewarding costly cooperation or punishing failures to cooperate that are significant for the persons being rewarded or punished but that cost the person administering them little or nothing. These include honor and esteem for the cooperator and dishonor and social isolation for the non-cooperator.¹⁵

The effect of all these possibilities is to convert a collective action problem into a coordination problem. However, there are still important differences between the coordination problem faced by predators and the coordination problem faced by those who merely want to resist predation. It will not always, or perhaps often, be true that the actions of individuals who cooperate in resistance will have a significant impact on the outcome, but individuals who cooperate in predation can always be given a share of the gains. Moreover, if resistance requires the cooperation of people who are part of an existing stable group of some sort, then any such groups that exist may be too small to resist a powerful organization of predators effectively or lack the means of mobilizing whatever conventions or social norms they have developed to enable resistance. Predators, however, can make use of soldiers of fortune, who can be readily attracted by the expected gains from predation.¹⁶

Moreover, the mechanisms that can enable resistance to predation can also be used to facilitate it. Every effective military organization knows how to make use of the mechanisms just described to make soldiers willing to risk their lives in combat, many successful predators have avoided conflicts within their ranks by using their own families as the core of their organizations, and many predatory organizations try to attract support and avoid resistance by claiming to serve collective interests or to be enforcers of valued group norms. All these devices not only make predatory organizations more effective but also make predation more profitable for their leaders, since the less leaders have to use the gains from predation to compensate their followers, the more of the gains they can keep for themselves. This helps explain why it is possible to debate whether the Crusades during the Middle Ages were motivated by religion or predation,

15. Since cooperation is important for human survival, there is reason to believe that humans have evolved psychological mechanisms that facilitate it. One of these may be an inborn inclination to punish noncooperation or violations of social norms (Bowles and Gintis 2004). For some experimental evidence in support of this possibility, see Fehr and Gächter 2002. For a discussion of the collective action problem from an evolutionary perspective, see Wilson 2002.

16. For a discussion of how rebels can overcome the collective action problem, see Lichbach 1995. See also McAdam, Tarrow, and Tilly 2001. For a discussion of coordination mechanisms, see Chwe 2001.

there is often disagreement about whether armed groups are bandits or revolutionaries, and there can now be disagreement about whether “greed or grievance” motivates civil wars in various parts of the world (Berdal and Malone 2000).

The Protection Business

In hunter-gatherer societies, the gains from economic predation are small, while the survival of the group requires cooperation among its members. Thus acts of predation by members of the society against other members are mainly carried out by individuals, and since the group is small they can be punished in informal ways. Moreover, the gains from collective predation by one such society against another are also likely to be small, and therefore conflicts among them are likely to be the result of acts of individual predation by members of one group against members of another. This may lead to collective retaliation and counter-retaliation, but such feuds eventually die out and are not a threat to the independence of the groups (Boehm 1987).

With the development first of settled agricultural communities and then of trade among them, however, the gains from organized predation increase. It can take three forms: raiding (or banditry), in which organized groups attack settled communities or traders and carry off their harvests or other goods; the capture of valuable agricultural land or trading routes by one group from another; and long-term exploitation, in which an organized group of warriors acquires control over settled communities and subsists off their produce.¹⁷

As already noted, groups engaged in economic predation are sometimes said to be in the protection business. This is not always just a bad joke. And in the case of settled predators, this expression conveys an important insight.¹⁸

The relation between settled economic predators and the people they exploit is similar to the relation between a parasite and its host or between a farmer or rancher and the domesticated animals that support him, in

17. This is the main theme of McNeill 1982. McNeill calls economic predators “macroparasites,” that is, “men who, by specializing in violence, are able to secure a living without themselves producing the food and other commodities they consume” (1982, vii).

18. Because groups of economic predators may have an interest in protecting the people they prey upon, they are now sometimes called “mafias.” The idea is developed in Tilly 1985. For an analysis of the Sicilian mafia that argues that it really is in the protection business, see Gambetta 1993. For a recent study of the so-called Russian mafia by one of Gambetta’s students, see Varese 2001. These ideas are developed further in Volkov 2000 and 2002. The seminal work on states as organizations in the protection business is Lane 1958. Lane’s idea is the main theme of Glete 2002. See also Levi 1988.

that settled economic predators have an interest in the productivity of the people who support them and in protecting them from other predators. This is why the “protection” that is “sold” by such predators can be the genuine article:

the king is . . . the owner of the country. Like the owner of a house, when the wiring is wrong, he fixes it.¹⁹

Thus even an economic predator, if he is engaged in the long-term exploitation of a settled community, would have an interest in providing some of the core services we associate with governments. But the terms of trade between protector and protected can vary enormously. At one extreme a protector may be a pure predator, the wealthiest person in the territory that he controls, and provide few benefits for anyone else. At the other extreme he may be hired by the people he protects, many of whom enjoy greater wealth and job security than he does (Grossman 2000). The extent of the services offered by a protector, and the terms on which they are provided, will be influenced by the relative bargaining power between the protector and his customers. And this will be influenced in turn by the amount of competition the protector faces and the ability of the people he protects to organize themselves in bargaining with him.

Consider, for example, the famous Japanese film by Akira Kurosawa called, in English, *The Seven Samurai*. In this film, peasants in a Japanese village are plagued by roving bandits who regularly steal their harvest, and they pool their resources to hire samurai warriors to protect them. They find seven rootless samurai warriors who agree to help and who succeed in defeating the bandits. In this story the exploiters are roving bandits, the villagers are able to organize because they are a small community, and the samurai warriors are rootless and unorganized. The story would end differently if the bandits were strong enough to establish control over multiple villages, which would be unable to coordinate resistance against them, or if the surviving samurai warriors remained in the village after the battle that ends the film and extorted much greater payments from the villagers. The last possibility exemplifies the fact that even groups that manage to organize themselves in defense of exploitation by others are exposed to the danger of being exploited by the people they have empowered to defend them.

Since the people a predator exploits provide the means of protection against competing predators, the possibility of competition from other predators may increase an established predator’s interest in their produc-

19. A statement by an Italian monarchist, quoted in Banfield 1958, 27. For an analytical treatment of this point, see Olson 1993, who calls settled economic predators “stationary bandits.”

tivity. And the fact that the people he exploits are not dumb animals but are capable of resisting efforts to mobilize them against competing predators may increase his interest in their well-being. Moreover, these two possibilities can reinforce each other, since (1) an effective means of resisting an established predator is to support a competitor, (2) competitors may have an interest in facilitating resistance to an established predator, and (3) measures that a predator takes to increase the productivity of his prey may also increase the ability of his prey to organize themselves to resist him. All these things may be true of both competing predators from within the territory controlled by an established predator and those outside it.

The Social Contract

A state is usually defined as an organization that has several properties, among which are (1) an organized ruling group (the government) that (2) can successfully use the threat of force to compel individuals (3) within a well-defined territory (4) to surrender economic resources (taxes), which are used to support not only the ruling group but also its (5) regulation of the population it controls and (6) defense of its control from other organized groups, both within the territory it controls and outside it. How are we to explain the development of organizations with these properties?²⁰

Max Weber wrote:

“Every state is founded on force,” said Trotsky at Brest-Litovsk. That is indeed right. If no social institutions existed which knew the use of violence, then the concept of “state” would be eliminated, and a condition would emerge that could be designated as “anarchy,” in the specific sense of this word. (1946, 78)²¹

It is clear that in this passage Weber has in mind not just any use of force or violence but the organized use of violence. To understand the development of states, then, we must explain the development of organizations capable of engaging in killing and destruction and then show why they would be interested in creating something that would have the properties ascribed to states. A plausible explanation can be found in the incentives

20. See, for example, the definitions of a state by Weber (1946) and Finer (1997, 1–94). Unfortunately, much of the literature on the rise of the state familiar to students of international politics focuses on the development of the European state in the context of feudalism. But human organizations with the defining properties of a state have long existed and have developed in many places. In addition to Finer, see Johnson and Earle 1987 for a survey of the anthropological literature.

21. Note that in this passage Weber equates anarchy with the *absence* of violence.

for economic predation, and defense against it, created by the development of settled agricultural communities and trade among them (North 1981).

But neither economic predators nor people who organize to defend themselves against them have an interest in violence for its own sake, any more than a man who holds up a liquor store wants to shoot the person behind the counter. What they want is to be able to profit from an agreement whose terms are determined by the threat of violence. It is the provisions of these agreements that determine the peculiar features that distinguish any particular organization that has the properties listed previously from all the others.

Any particular state, therefore, can, like a business firm, be understood as a network of contracts (Aoki, Gustafsson, and Williamson 1990). One set of contracts regulates the internal functioning of the ruling group (the government), another regulates the relation between the government and the population whose behavior it tries to control, and a third regulates the relation between a government of one territorial area and the governments of other territorial areas (Glete 2002, 1–41). Unlike the contracts that organize a business firm, however, the disagreement outcome in negotiating the contracts that define a state is determined by contests in violence. Moreover, these contracts are never entirely independent of each other, and the enforcement of all of them is problematic.

Weber famously defined a state as

a human community that (successfully) claims the *monopoly of the legitimate use of physical force* within a given territory. (1946, 78; emphasis in original)

It is the contractual basis for the functioning of a state that explains the role of legitimacy in its organization.

The idea of a fair bargain can be a way of thinking about the properties of an ideal or just state, and an actual state that did not satisfy those properties might be considered illegitimate because it was unjust. But even the relation between master and slave is subject to bargaining, and so also are unjust states governed by tacit or explicit agreements. Even though the bargaining power of slaves may be insufficient to compel the abandonment of slavery, it can be enough to punish a master who has deviated from rules governing his behavior that were the product of earlier bargaining between them, and such rules, however vague or implicit, might be said to define what is and what is not legitimate behavior for a master or for a ruler.

Just as the terms of any contract reflect the relative bargaining power of the people who sign it, so standards of legitimacy understood in this

way reflect the relative bargaining power of the component parts of a state. A change in relative bargaining power will therefore make it possible to renegotiate the contract. But, as in the relation between labor and management, one of the determinants of relative bargaining power is the relative ability of the participants to organize, which is affected in turn by their ability to coordinate their actions. Thus the mere availability of the idea of a radically different type of contract, if it becomes common knowledge, can change the relative bargaining power of ruler and ruled, since such an idea can facilitate the organization of resistance to the government defined by the current contract. The relation between “realistic” standards of legitimacy and “utopian” ones is therefore complex, and the way people talk about government can undermine it (Mannheim 1936).

Anarchy and Hierarchy Reconsidered

Kenneth Waltz wrote:

The parts of domestic political systems stand in relations of super- and subordination. Some are entitled to command; others are required to obey. Domestic systems are centralized and hierarchic. The parts of international-political systems stand in relations of coordination. Formally, each is the equal of all the others. None is entitled to command; none is required to obey. International systems are decentralized and anarchic. The ordering principles of the two structures are distinctly different, indeed, contrary to each other. (1979, 88)

The distinction described in this passage is the foundation of what came to be known as “structural Realism.” But it should now be clear that it rests on a confusion, a confusion that is fostered by the words that Waltz uses to describe the distinction he wants to make. The opposite of a centralized system is not anarchy but a decentralized one. *Anarchy* is *an-* plus *-archy*. It refers to the absence of a leader or ruler, which structural Realists equate with the absence of any institutional structure.²²

Thus this passage conflates two different distinctions. One is between a centralized and a decentralized institutional structure, and the other is

22. Strictly speaking, the opposite of anarchy is not hierarchy but “archy,” a word that does not exist by itself in English but must be qualified by some prefix that describes what sort of “archy” it is (e.g., monarchy or oligarchy).

between relations among people that are governed by an institutional structure and those that are not.²³ Structural Realism therefore begs the question of whether peace requires some sort of *-archy* or could be achieved instead by a decentralized institutional structure.²⁴

This confusion has been compounded by the confusion between the absence of government and Hobbes's state of nature. If a world without governments is a world without states as commonly defined, then people who live in a world of hunter-gatherer societies might be said to be in the state of nature, but it would not be the state of nature described by Hobbes.²⁵ Hobbes's state of nature was a world of competing economic predators, any of whom could become the nucleus around which what Hobbes called a commonwealth could be constructed and some of whom did. And there is nothing in what Hobbes wrote that implies that a world of commonwealths must have the properties that Hobbes attributed to what he called the state of nature.

Like Hobbes, many students of international politics do not distinguish between individual acts of predation and predation by organized groups. This failure helps support the view that the difference between government and anarchy is that under government there is someone to enforce contracts and property rights and in anarchy there is not, a view that makes plausible the use of the Prisoner's Dilemma game as a way of describing anarchy or the state of nature. A settled economic predator with sufficient interest in the productivity of his subjects may be interested in resolving conflicts among them and defining and enforcing their property rights, but there is no enforcer of the contract (implicit or explicit) that governs the relations between the predator and his subjects, any more than there is an enforcer of any contracts he might have made with economic predators in other territories that delineate which territories they each control. In both cases, as in bargaining between master and slave, any partic-

23. To appreciate the difference, think about the distinction between having a commonly known "rule of the road" and not having one. A failure to appreciate this distinction is what distinguishes structural Realism from what is commonly called the "English school" of writers about international politics. For a representative example of the English school, see Bull 1977. For a recent discussion of the English school, see Keene 2002.

24. Hans Morgenthau, it should be noted, distinguished clearly between a system without any institutional order and a system with a decentralized institutional structure and argued that the fundamental property of modern international politics was not the absence of effective international norms or law but its decentralized institutional structure (1948, part 6). See also the discussion of this issue by Martin Wight, a founding member of the English school, in Wight 2002, chaps. 9 and 10.

25. See the analysis of the anthropological and archaeological data in Kelly 2000. It was perhaps in part the European experience of hunter-gatherer societies in the New World that motivated Rousseau's criticisms of Hobbes (Meek 1976).

ular agreement is enforced only by threats to repeat the forceful bargaining that produced it.

Because a ruler commands an organization, and violators of his edicts do not, a ruler will be able to confront violators with take-it-or-leave-it demands. Even if there are many violators, all of whom meet with a violent response, this violence will not count as warfare, because it will not consist of a military contest between organized groups. It is the ruler's monopoly of the *organized* use of force that explains his ability to enforce property rights without war, but his monopoly of the organized use of force exists only because the members of his organization cooperate in applying it, and any resistance to it is not organized. Thus the superior force of the state is not the cause of the reliability of agreements that are accepted in lieu of violence but its result, and when the agreements that support it unravel then so does the state. A potential cause of such unraveling is an attempt by the state to enforce too much.

Waltz's distinction between hierarchy and anarchy derives additional plausibility from the fact that the history of international politics is usually told as the history of warfare between or among independent states. This fosters the view that states exist independently of other states and some way needs to be found to prevent them from fighting each other. But every war ends in a peace settlement of some sort, and the states that participate in any given war were all products of some prior peace settlement. As Robert Randle said:

It is . . . wars and their settlements that have structured the state system of the modern era: they have provided the matrix for interstate relations, including the context of subsequent wars and their settlements. . . . Peace settlements . . . created the modern state system; they have characterized the relations of states and the international law of those relations; and it is through them, in part, that the modern state became what it is. (1973, 506)

Thus sovereignty does not reflect an absence of agreements but is itself the result of agreement. Indeed, Finer lists, as one of the defining properties of a state, the fact that it is

recognized by other similarly constituted states as independent in its action on its territorially defined . . . population, that is, on its subjects. This recognition constitutes what we would today call its international "sovereignty." (1997, 2–3)

There is no external enforcer of the agreements constituting a state's sovereignty, but neither is there an external enforcer of any of the other agreements that constitute a state.

A history of modern international politics told as the history of peace settlements would be a history that revealed the institutional development of the European state system, which was eventually extended to encompass the globe.²⁶ It would show that both the interstate system and the states that make it up are constantly being renegotiated and that the modern state is as much the product of agreements among states as it is of agreements between governments and the populations they govern. When states use force to renegotiate a previous peace settlement they appear to be the source of the problem, but when a new agreement is negotiated they reemerge as part of the solution. And no valid argument has been given that shows that they could not be part of a long-lasting peace settlement.

Constructivists, in criticizing structural Realism, emphasize the fact that the state is a social construction whose origins lie as much in the international system as in the societies they govern. However, the alternative they offer to simply taking states as given, as structural Realism does, is an analogy with the process by which individuals acquire their identities through socialization. But individuals do not negotiate their identities by the use of force. States are the product of a process by which groups of individuals with well-defined identities use violence to bargain over the institutional structures that will regulate conflicts among them. And one of the issues to be bargained over may be the nature of the culture that they will subsequently be part of.

The Global Constitution

As we have seen, Waltz characterized the interstate system as “decentralized.” If this does not mean simply the absence of an institutional structure, as the word *anarchy* implies, what does it mean?

In thinking about the answer to this question, we must distinguish between a world of states with an uncontested institutional structure and a world in which the structure is contested. This is something we are accustomed to doing in talking about the internal structures of states. The institutional structure of the U.S. government once was the subject of a violent contest, but there seems little prospect of that happening again in the immediate future. Thus the institutional structure of the U.S. now conveys a great deal of useful information about how life in the United States is conducted. This could not be said of contemporary Colombia, however, or Afghanistan or the Balkans. We can imagine possible institutional structures for the people who live within those areas and think about what sort of institutions, if any, might prove to be acceptable to everyone within them who might be in a position to use violence to contest them. If we are

26. See, for example, Holsti 1991; Osiander 1994; and Ikenberry 2001.

to think usefully about the possibility of a peaceful global order we must do the same for the state system itself.

A global order that consists of a state system, rather than a world government, would have an institutional structure that did not include an organization at the global level with the defining properties of a state listed previously. Thus there would be no global organization that could use threats of force to tax individuals in the constituent states. And, therefore, since the states themselves would constitute the institutional structure of the global order, the state system would be defined by the three sets of agreements, listed earlier, that defined the constituent states.²⁷

But if a world of sovereign states is simply a world without a global government, then agreements defining the external sovereignty of states might include provisions that regulate the other two sets of contracts that define the constituent states. And, indeed, this has always been true. Wars were fought in the eighteenth century to determine the ruling family of one or another member of the European state system; the nature of the government of France was one of the issues determined by the settlement that ended the wars of Napoleon; and the nature of the states that will govern the Balkans, as well as their territorial boundaries, is still being negotiated between the members of NATO and the people who live in that area.²⁸

If a state is unable to control individual or small-group acts of predation within its territory, it will be unable to prevent people who inhabit its territory from engaging in acts of predation on the territory of other states, and therefore a minimal condition for the external sovereignty of states is that they be able to exercise internal sovereignty. Thus over time the state system invented by the Europeans has come to resemble the scheme for regulating conflicts between or among ethnic groups described by Fearon and Laitin in their study of ethnic conflict: governments are responsible for preventing individuals inhabiting their territory from engaging in acts of predation against people who inhabit the territory of other governments, a division of responsibility that diminishes the number of occasions for violent conflicts between or among states (Fearon and Laitin 1996). The current so-called war on terrorism is based on this principle, but it is

27. See the recent discussion of the global constitutional order in Bobbitt 2002, especially book II.

28. Contrast this discussion with Krasner's (1999) discussion of sovereignty. Krasner claims that, "According to the Westphalian model relations between rulers and ruled ought not to be subject to any external actors" (73). The many exceptions to this supposed norm lead Krasner to characterize the "norm" of sovereignty as "organized hypocrisy." Note that it is important to distinguish among (1) any conditions that are attached to the interstate agreements defining states' sovereignty, (2) interventions to enforce those conditions (e.g., the enforcement of basic human rights), and (3) the use of force to renegotiate those conditions.

merely the latest installment of a long story.²⁹ Fearon and Laitin list a number of properties of ethnic groups that may give them a comparative advantage over other groups in controlling acts of predation by their members. A similar claim might be made for states in comparison with other possible enforcement agencies.

Peace and the State

The difference between the existence of a monopoly of the legitimate use of force within states and the absence of such a monopoly among them is what Waltz really had in mind when he distinguished between hierarchy and anarchy. And a side effect of the existence of multiple states, of course, is that they can use their ability to support organized military forces to engage in violent conflicts with each other, which they frequently do. Since war requires the organized use of force, one might think that a world of states would necessarily make war within states impossible, while war among them could not be ruled out. This is why structural Realism has seemed plausible to so many people.

But a monopoly of the legitimate use of force can be lost, and therefore wars within states are not impossible. And the mere possibility of war among states does not imply that it will occur with any significant probability. Thus the inference from the institutional structure of a state system to the incidence of war within it is not valid.

It is true that, for war to occur within states whose institutional structure is not already contested, an organizational problem has to be solved that need not necessarily be solved for war to occur between states. If war is to occur within the territory of an existing state that enjoys a monopoly of the legitimate use of force, then the state's military forces must be divided and/or a new military force must be created to oppose the one previously controlled by the government.³⁰ States with a monopoly of the legitimate use of force within their territories, however, may maintain armies ready and able to fight each other on short notice.³¹

These differences imply that there might be an institutional impediment to war within some states that does not exist between some states.³² However, there are other, less obvious implications of these differences

29. Much of the story is told in Thomson 1994.

30. For a description of this process in conjunction with the U.S. Civil War, see Bensel 1990.

31. There may, however, be a need to organize an alliance if a war is to be fought among states.

32. But not all states. Canada, for example, would have some organizational problems to solve before it would be prepared to fight a war against the United States.

whose effects are contrary to this obvious one. It is possible, for example, for groups of individuals within states who contemplate violence to coordinate their expectations sufficiently that they each expect to profit from using force, yet they lack an organization that could commit them all to an agreement that they might all prefer to the expected consequences of using force. In that case peace within the territory of an established state might be harder to achieve than peace between states.

Consider, for example, the Los Angeles riot in the aftermath of the acquittal of the Los Angeles policemen who had been videotaped beating a black man, Rodney King. Shared outrage at the verdict in that trial caused many blacks to congregate around the same traffic intersection, and their observation that the police chose to withdraw from the scene rather than try to control the crowd told them that individuals who decided to use force would not face either effective opposition or especially dire consequences. Thus each was free to vent his or her rage against white people, or appropriate property from the many stores in the area, until a large enough military force was organized to oppose them. But if the leaders of this military force contemplated negotiations with the rioters there would be no one to negotiate with, and therefore quelling the riot required sufficient use of force to demonstrate to all the individuals involved that the balance of power between them and the police had been reversed.

Spontaneous demonstrations such as this one can have immediate revolutionary consequences if they occur in the capital of a centralized state (e.g., Paris), and if the public authorities are unable to alter the expectations that support them they can lead to recurring violence over the long run.³³ However, if no single organization develops that can negotiate an agreement and then persuade the dissidents to accept it, then a negotiated settlement of such conflicts may be impossible. For example, in attempting to negotiate a settlement between rebels in Kosovo and the government of Serbia prior to the military conflict over Kosovo, the U.S. government had difficulty in finding someone who could reliably speak for the rebels. And one of the main incentives for the government of Israel to agree to the creation of a Palestinian state seems to be the possibility that an agreement negotiated with such a state might reduce decentralized violence by Palestinians against the citizens of Israel.³⁴

33. As a former adviser to a Chinese leader has been quoted as saying, "There are so many people with grievances. They'll wait for some public signal, and then they'll come together when they know others will do the same thing" (Ziegler 1997, 20).

34. For an argument that a lasting peace between Israel and the Palestinians requires one state and not two, see Said 1999.

What Next?

Thus not only can states with a monopoly of the legitimate use of force within their territories make war with other states; they can make peace with them as well, which they also do. This is something that organizations without a monopoly of the legitimate use of force are unable to do. The institutional structure of a state system does not tell us why the peace that states make among themselves could not be as lasting as the peace some states, but not all, have made within their territories.

Moreover, precisely because of the institutional impediment to the renewal of conflict that a government would entail, the members of warring groups may have less confidence in the terms of a peace settlement that creates a common government than one that provides for separate states with separate military forces. Thus even if there is a subsequent conflict between or among the resulting states, the extent of the violence may be less than if an agreement had not been reached, and the existence of separate states may not be its cause but rather may reflect prior expectations that conflict was likely.

To understand the recurrence of war in a system of states, we must therefore understand why states that could make peace with each other make war instead. Moreover, every war takes place in a world that was created by some prior peace settlement and will end with another one. So to understand the recurrence of war, we must explain why peace settlements do not last. To do that we must look more closely at the relation between bargaining and war.³⁵

35. For two seminal articles on this subject, see Fearon 1995b and Powell 1996. Much of what I have to say in the following chapters is based on ideas developed in those articles.

Bargaining and War

Kenneth Waltz's "third image" of the causes of war (1959), which was the foundation for what came to be known as "structural Realism" or "Neo-realism," was inspired, as we have seen, by Jean-Jacques Rousseau's description of a world of predatory rulers. But it is unclear from what Rousseau wrote why a world of predatory rulers had to be as conflictual as it was, since, as I pointed out in the previous chapter, competing predators would appear to have an incentive to reach agreements to share the benefits of rule among themselves.

Kant, like Rousseau, thought it was obvious that a world of predatory rulers would be a world in which war was frequent, but, bad as this was, he believed it was nonetheless better than the alternative, since a lasting peace among predatory rulers would have prevented achievement of the justice and prosperity that he expected would be the eventual consequence of recurring wars. Once justice and prosperity had been achieved, he thought, peace might be possible. But like Rousseau and nearly all other writers on this subject, he had little to say about why war occurs at all.

Thus if we are to evaluate these ideas we must think about why wars occur, and we should begin by thinking about wars among predatory rulers. To do that, we must write down what seem to be the relevant properties of such a world and see if they have any clear-cut implications for the occurrence of wars.

Warring Predators

As we saw in the previous chapter, a contest in forcible disarmament is not the only form that a contest in killing and destruction might take, but it is the obvious place to begin in thinking about wars among competing economic predators. There are two reasons why one economic predator might expect to gain from forcibly disarming another: if one controls valuable territory, then the other might expect to gain from capturing it; and if both are trying to exploit the same producers, then either could increase his gains by eliminating the other.

A contest in forcible disarmament might lead to the disarmament of either side, and the probability with which either outcome might occur would be a function of the military capabilities of both sides. Thus such a contest resembles in some ways an athletic contest—though it would be more accurate to say that many athletic contests and other games were designed to resemble military contests. This helps account for much commonsense reasoning about war, which is based on the idea that wars are contests that either side might win or lose, with a probability that is determined by their relative power or military capabilities.

However, the analogy between contests in forcible disarmament and athletic contests might lead one to ask whether wars can have only two outcomes, since athletic contests can end in ties, and it is often said of a war that it ended in a stalemate. But athletic contests end in ties because the rules by which they are conducted specify when the game ends and the score might be tied at that point. When people want to avoid ties, then the game is continued until one side or the other wins. There are no rules that specify when a war should end, and therefore if a contest in forcible disarmament ends before either side has been disarmed it is because the combatants chose to end it—which they might have done because they saw no immediate prospect of either defeating the other.

Thus the problem with much commonsense reasoning about war is not that it assumes that wars have only two outcomes but that it overlooks the fact that such contests can be interrupted if the combatants choose to stop fighting, and therefore it assumes that after war begins states no longer face a choice between fighting and not fighting. But if rulers can decide to stop fighting or continue, they can also make any decision to stop fighting conditional on the acceptance of an agreement of some sort. Economic predators, for example, could agree on a redivision of the valuable territory that they are fighting over, instead of continuing to fight until one or the other had been disarmed. And, indeed, many wars have ended in just this way. But this is something they could have done without fighting at all, and therefore the fundamental problem in explaining the occurrence of war is to explain why the participants had to fight before reaching an agreement that settles whatever is in dispute between them.¹

This implies in turn that, even if there are only two ways that a contest in forcible disarmament can end, there can be many possible outcomes of a war, since a war can be ended by an agreement, and there are many possible agreements that might be reached. In a contest between economic predators, for example, the territory they control could in principle be divided in indefinitely many ways. And therefore, if we are to explain why states fight on the basis of their expectations about the likely consequences

1. This is the main theme of Blainey 1988.

of fighting, we must take into account their expectations not only about the likely outcome of a contest in forcible disarmament but also about the outcome of the bargaining process that might accompany it.

The first major writer to point this out and attempt to determine its implications was a Prussian military officer, Carl von Clausewitz, who lived from 1780 to 1831 (Clausewitz 1976). Like Hobbes, Clausewitz wrote in an arresting style that lends itself to quotations taken out of context. In addition, he never finished his great treatise, *On War*, and it was published by his wife after his death. As a result, he has been misunderstood almost as often as he has been quoted. Moreover, while his analysis was surprisingly modern and sophisticated, we can now see that at the heart of it is what is commonly called the bargaining problem, whose full complexity has only become apparent as a result of the analytical techniques developed by game theorists.

Clausewitz wrote, “War is . . . an act of force to compel our enemy to do our will” (1976, 75). From this it followed, he claimed, that “the aim of warfare is to disarm the enemy,” since “[i]f the enemy is to be coerced you must put him in a situation that is even more unpleasant than the sacrifice you call on him to make,” and “[t]he worst of all conditions in which a belligerent can find himself is to be utterly defenseless” (77). “Force,” he wrote,

is thus the *means* of war; to impose our will on the enemy is its object. To secure that object we must render the enemy powerless; and that, in theory, is the true aim of warfare. (75; emphasis in original)

But the enemy can be expected to resist this outcome, and this resistance must be countered if he is to be disarmed. “Each side, therefore, compels its opponent to follow suit; a reciprocal action is started which must lead, in theory, to extremes” (77).

Statements such as these have led some people to interpret Clausewitz as an apostle of total war. But such an interpretation overlooks the qualifying phrase *in theory* that appears in these quotations. In practice, Clausewitz wrote, war does not usually look like that at all.

In practice, Clausewitz wrote, “war is simply a continuation of political intercourse, with the addition of other means,” a statement that is often quoted but, in light of such statements as the ones quoted previously, often interpreted as mere cynicism. However, Clausewitz meant this statement to be taken literally:

We deliberately use the phrase “with the addition of other means” because we . . . want to make clear that war in itself does not sus-

pend political intercourse or change it into something entirely different. (1976, 605)

Thus at the heart of Clausewitz's discussion of war in practice, or, as he sometimes called it, "real war," is the fact that war is typically accompanied by the same bargaining process that preceded it and that will continue after it ends.² And the reason this is possible is that, as Clausewitz put it, "war does not consist of a single short blow," and therefore negotiations with the enemy need not await his complete defeat (79).

One implication of this fact, Clausewitz wrote, is that "real war" may actually consist not of a contest in forcible disarmament that is interrupted by a negotiated settlement but of a contest in killing and destruction in which the adversaries do not even try to disarm each other. Rulers may instead simply fight over a particular piece of territory or even engage in military operations whose object "is neither to conquer the enemy country nor to destroy its army, but simply *to cause general damage*" (Clausewitz 1976, 93; emphasis in original). "What is more," he wrote,

a review of actual cases shows a whole category of wars in which the very idea of *defeating the enemy* is unreal: those in which the enemy is substantially the stronger power. (91; emphasis in original)

Thus Clausewitz claimed that a ruler could be optimistic about the outcome of war, even though he was not optimistic about defeating the enemy in a contest in forcible disarmament—a possibility that is overlooked entirely by most commonsense reasoning about war.

"Warfare thus eludes the strict theoretical requirement that extremes of force be applied," Clausewitz wrote, and "[t]he probabilities of real life replace the extreme and the absolute required by theory."

Once the extreme is no longer feared or aimed at, it becomes a matter of judgment what degree of effort should be made; and this can only be based on the phenomena of the real world and the *laws of probability*. . . . reality supplies the data from which we can deduce the unknown that lies ahead.

From the enemy's character, from his institutions, the state of his affairs and his general situation, each side, using the *laws of probability*, forms an estimate of its opponent's likely course and acts accordingly. (80; emphasis in original)

2. Clausewitz sometimes calls war in theory "absolute war," and he sometimes refers to "real wars" as wars with "limited aims" (1976, book 8).

However, while “[t]heory must concede all this,”

it has the duty to give priority to the absolute form of war and to make that form a general point of reference, so that he who wants to learn from theory becomes accustomed to keeping that point in view constantly, to measuring all his hopes and fears by it, and to approximating it *when he can* or *when he must*.

A principle that underlies our thoughts and actions will undoubtedly lend them a certain tone and character, though the immediate causes of our action may have different origins, just as the tone a painter gives to his canvas is determined by the color of the underpainting. (581; emphasis in original)

What Clausewitz seems to be saying is that, while states that are fighting may not actually try to disarm each other, they must bear in mind the fact that they could, and absolute war, even though it never occurs, must be the “measure of all their hopes and fears.”

But “[i]f theory can effectively do this today,” he wrote,

it is because of our recent wars. Without the cautionary examples of the destructive power of war unleashed [by Napoleon], theory would preach to deaf ears. No one would have believed possible what has now been experienced by all. (581)

It is striking to compare this statement with Thomas Schelling’s comment about the limited nature of the Korean War: “It is a strange spectacle, and indeed what makes it plausible is only that it actually occurred” (1960, 130). The expectations of Clausewitz’s readers were conditioned by experience of the limited wars of the eighteenth century. The expectations of Schelling’s readers were conditioned by experience of the total wars of the twentieth century. But Clausewitz and Schelling agree that, as Schelling put it, “[w]ar is always a bargaining process” (142), that the nature of wars is determined by states’ choices rather than the technology that is available, and that to explain why they choose to fight the wars they fight one must understand the bargaining process that wars are part of.

But this means that there are two fundamental puzzles about war and not just one: we must explain not only why states must fight before reaching an agreement, when they could have reached an agreement without fighting, but also why they chose to agree to fight only a limited war, when the outcome of a contest in disarmament would have been different.

To someone familiar with the modern literature on bargaining, Clausewitz’s solution to both puzzles practically leaps off the page. It has two parts. Here is the first:

if one side cannot completely disarm the other, the desire for peace on either side will rise and fall with the probability of further successes and the amount of effort these would require. If such incentives were of equal strength on both sides, the two would resolve their political disputes by meeting half way. If the incentive grows on one side, it should diminish on the other. Peace will result so long as their sum total is sufficient—though the side that feels the lesser urge for peace will naturally get the better bargain. (Clausewitz 1976, 92)

Translated into modern terminology, this says that a contest in disarmament (absolute war) is the disagreement outcome in any bargaining over the terms of a settlement that might substitute for war. Thus the more optimistic a ruler is about the outcome of absolute war, the better the terms he will demand and expect in any agreement he might accept instead, and vice versa; and if the demands of the two adversaries are compatible, an agreement can be reached without fighting.

Here is the second part of Clausewitz's solution to these puzzles:

When we attack the enemy, it is one thing if we mean our first operation to be followed by others until all resistance has been broken; it is quite another if our aim is only to obtain a single victory, in order to make the enemy insecure, to impress our greater strength upon him, and to give him doubts about his future. If that is the extent of our aim, we will employ no more strength than is absolutely necessary. (92)

This second statement says that if military operations are not designed to disarm the enemy, their purpose is to influence his expectations about what the outcome of absolute war would be, were it to be fought. Thus the function of "real wars" is to reveal information about the adversaries' military capabilities.³

Taken together, these two ideas raise two important questions. The first is whether, if rulers' expectations about the outcome of absolute war are sufficiently consistent, they would always be able to reach an agreement without fighting.⁴ The second is whether, even if this is not true, they might nonetheless only need to fight wars that are not very costly or even, perhaps, engage in other types of conflicts that, while inefficient, are nonetheless much less costly than military conflicts would be—interruptions of trade, for example.

3. This is the central theme of Blainey 1988. The idea is developed in Wagner 2000.

4. This is the main claim made by Geoffrey Blainey (1988).

Whatever the answers to those questions may prove to be, Clausewitz's two ideas clearly imply that the belief that the increasing costliness of war might in itself be sufficient to make war obsolete is unwarranted: the costliness of absolute war might make an agreement to avoid it desirable, but rulers can nonetheless choose to fight wars that they expect to be less costly instead. That is why the belief that World War I had demonstrated that war was too costly to be repeated was misguided and may even have contributed to the occurrence of World War II.

Let's Make a Deal

Clausewitz's analysis of war gives us further reason to take seriously Kenneth Waltz's analogy between wars and strikes.⁵ It implies that, to understand what happens on the battlefield and its consequences, we must understand not only the military contest but also the bargaining process that accompanies it. This poses a very complicated set of problems, and thus we should not be surprised if formal models prove to be necessary in thinking about it.

Let us begin by asking whether we should expect predatory rulers with consistent expectations about the outcome of a military contest to be willing to reach a peaceful agreement dividing valuable territory between them rather than fight over it. A contest in forcible disarmament (Clausewitz's absolute war) resembles a costly lottery, since there is some probability that either side might win. Winning such a contest would imply control over all of the territory in dispute. By agreeing to divide it up rather than fight, however, rulers could avoid both the costs and the risks associated with a military contest. Thus the choice between a war and a negotiated settlement involves a choice between a sure thing and an uncertain prospect.

Such a choice is represented in figure 5. Arrayed along the vertical axis are all the probabilities of winning the contest, from zero to one. Arrayed along the horizontal axis are all the possible fractions of the territory in dispute that a ruler might receive, from zero to one. The lines in the figure represent possible points of indifference for some particular ruler between getting some fraction of the territory for certain and fighting a contest for all of it with some specific probability of winning. The straight line, for example, represents the preferences of a ruler who is always indifferent between getting some fraction q of the territory and fighting a contest in which he expected to get all of it with a probability p of the same size. The curved line, on the other hand, represents the preferences of a ruler who, if

5. See the discussion at the beginning of the previous chapter.




Fig. 5. Choice between a lottery and a sure thing

confronted with some probability p of winning all the territory, would accept a lesser fraction q for certain instead. (In the figure, for example, this ruler would accept 60 percent of the territory as equivalent to the value of a contest for all of it that he had a probability of winning of .8.)

We have seen that Clausewitz was not bothered by the idea that expectations about the outcome of war could be represented by probabilities. Nonetheless, it is important to be clear about what these probabilities represent. The modern answer is that they are personal or subjective probabilities, which means they represent the odds at which a ruler would be prepared to bet on the outcome of a war. Thus they represent points of indifference between the gamble associated with war and a hypothetical lottery with known probabilities leading to the same outcomes, and therefore they are really just preferences. That does not imply that they are arbitrary, however. Rather, they incorporate all the information that a decision maker believes to be relevant to determining the outcome, in the same way that a person who bets on the outcome of a sporting event tries to take into account everything that he or she knows about the contestants and believes to be relevant.

Thus figure 5 is just a way of summarizing the preferences of someone choosing between a sure thing and a lottery: the probabilities represent points of indifference between the actual lottery and some hypothetical lottery with known probabilities, and the lines in the figure represent points of indifference between this lottery and possible divisions of the prize. Moreover, there is no right answer to the question of what either

should be. The two lines in the figure merely represent two possible sets of preferences—there are indefinitely many possible lines like the one that is bowed upward in figure 5 and indefinitely many that might sag downward as well. S-shaped curves, or curves with more complex shapes, are also possible. (You might ask yourself what your points of indifference would be, if the quantity at stake were a sum of money and the probabilities were actual gambles.) Moreover, nothing I have said so far requires that any decision maker actually thinks in these terms at all: a decision maker need not map out how he would respond to all the possible choices he might confront in order to choose between some particular contest and some particular proposed compromise. Thus figure 5 helps us organize our thinking but does not necessarily represent the way a decision maker organizes his.⁶

While there may be many divisions of disputed territory that an individual ruler would prefer to fighting a contest for all of it, if a contest is to be avoided both rulers have to accept the same settlement, and making a settlement more attractive for one requires making it less attractive for the other, as figure 6 illustrates. Now each point on the horizontal axis represents a possible division of the territory in dispute, between a fraction that goes to the ruler on the left (q) and the remaining fraction that goes to the ruler on the right ($1 - q$). The left vertical axis represents the left ruler's probability of winning the military contest (p), and the right vertical axis represents the right ruler's probability of winning ($1 - p$). The curve starting at the left-hand side of the horizontal axis is the indifference curve from figure 5, and the other curve is the corresponding indifference curve for the other ruler. In figure 6, it is assumed that the probability that the ruler on the left will win is .8. The question we are interested in is whether there must be some division of the territory along the horizontal axis that both will prefer to fighting over all of it.⁷

6. Figure 5 illustrates the fact that we can use a divisible good to measure what a gamble is worth to someone, or a gamble to measure what the good is worth, but there is nothing that measures both independently of each other. If we use the gambles on the vertical axis to measure the value of various quantities of the good on the horizontal axis, then the probabilities on the vertical axis are von Neumann-Morgenstern utilities, which are the basis for contemporary expected utility theory. (The justification for the idea that people would want to maximize their expected utility is that, if the value of the outcome of every choice is measured by the probability of winning the same gamble, then maximizing expected utility is equivalent to maximizing the probability of winning that gamble.) Thus there is little connection between "utility" as defined by expected utility theory and the classical concept of utility, which presupposes a way of measuring levels of personal well-being. For a useful introduction to expected utility theory, see Raiffa 1968.

7. For an influential discussion of this question, see Fearon 1995b. For some criticisms of Fearon's answer, see O'Neill 2001.




Fig. 6. Choice between a contest and a bargain

As figure 6 is drawn, there are many such divisions: the ruler on the left will prefer any division giving him more than 60 percent of the territory, and the ruler on the right will prefer any division giving him more than 10 percent. Thus any division giving the ruler on the left between 60 percent and 90 percent of the territory will be preferred by both rulers to fighting a contest for all of it. Moreover, it is easy to see from figure 6 that this will be true regardless of how likely it is that the ruler on the left will win: reducing p will shift the range of possible agreements to the left along the horizontal axis (as Clausewitz claimed it would), but it will always exist.⁸

The reason this is true is that the lines of indifference between contests and bargains represented in figure 6 are bowed upward, which means that, for every probability of winning, both rulers would accept a smaller fraction of the territory in dispute as a substitute for fighting a contest for all of it. If instead each would accept only a fraction equal to the probability of winning, there would be no agreement both would prefer to fighting and only one agreement they both would be willing to accept: one in which $q = p$ and $1 - q = 1 - p$. And if both curves sagged downward, there would be no agreement they both would accept as an alternative to fighting. Thus

8. This is why Wittman (1979) argued that the balance of power affects only the terms of a negotiated settlement that might be accepted in lieu of war but not whether war occurs or not. However, as we will see, figure 6 leaves unclear why war occurs at all, and therefore Wittman's reasoning was incomplete.

the answer to the question we started with is that there may be agreements both rulers would prefer to fighting, but there need not be.⁹

There are, however, reasons to believe that often there will be such agreements. One reason is that in choosing between sure things and gambles people often do have preferences that resemble the ones in figure 6. (Ask yourself the following question: If you had a lottery ticket that gave you a 50 percent chance at winning \$1,000,000, would you refuse to sell it if the most you could get for it was \$500,000?) And the other reason is that a war is not just a gamble; it is a very costly contest.¹⁰ Thus it may well often be true that a ruler confronted with the prospect of a costly and risky contest for valuable territory would be willing to accept a division of the territory giving him a fraction of it that is smaller than his probability of winning all of it.¹¹

A reason for thinking that this might not be true of both rulers is that, if the issue is a possible redistribution of territory that is already distributed between them, then any compromise agreement would entail one side's surrendering some territory to the other. If, for example, the ruler on the left controlled 50 percent of the territory in dispute but could defeat the other with a probability of .8, then he would prefer war to the status quo, and to avoid war the ruler on the right would have to appease him by surrendering some of his territory.¹² Often people seem to be willing to accept

9. By convention, indifference curves bowed upward like the ones in figure 6 are said to represent aversion to risk, curves that sag downward are said to represent risk acceptance, and indifference curves that are straight like the one in figure 5 are said to portray risk neutrality. Because the probabilities in these figures are also von Neumann-Morgenstern utilities, such curves are sometimes also said to represent diminishing marginal utility, increasing marginal utility, and linear utilities, respectively. All these terms are very misleading. These curves merely reflect an individual's points of indifference between gambles and sure things, and since such indifference points will be influenced by both the risk involved and the values an individual places on the objects in question, there is no way to know what actually determines them. It is best to think of them as simply reflecting an individual's preferences, like any other indifference curve.

10. Only the probabilities of winning are represented explicitly in figure 6, but not the expected costs of fighting. The expected costs would nonetheless affect the shapes of the indifference curves.

11. If this were not true, it would be hard to explain why wars are often ended by negotiated settlements before either side has been completely disarmed.

12. This example illustrates a flaw in commonsense reasoning about war that is more fundamental than the fact that it overlooks the possibility of negotiated settlements: even if a compromise is not possible, whether a ruler prefers war to the status quo or not depends not just on how optimistic he is about the outcome of war but on the status quo distribution as well. Even with a probability of winning of .8, the ruler on the left will prefer the status quo to war as long as he already controls at least 60 percent of the territory. Since compromises that are preferred to war may not always exist, this is something that we must bear in mind.

greater risks to avoid what they consider to be losses than they would accept to achieve possible gains of the same size, so the ruler on the right might have an indifference curve that sagged downward rather than the one attributed to him in figure 6.¹³

However, war is not just risky; it is also costly. Moreover, there is good reason to expect that the indifference curve of the ruler on the left would be bowed upward. Thus there might still be compromise settlements that both would prefer to fighting.

What have we learned from all this? When force is used not to disarm an adversary but to harm people or destroy their property (or, as Clausewitz said, “to cause general damage”), it is obvious that its purpose must be to compel an agreement that both the perpetrator and the victim would prefer to a continuation of the conflict, and therefore contests in punishment must be part of a bargaining process. A contest in disarmament, however, is a contest to determine how much punishment two adversaries can subsequently inflict on each other: the winner of such a contest can use force to punish the other without organized resistance.¹⁴ As Clausewitz wrote:

War is nothing but a duel on a larger scale. Countless duels go to make up a war, but a picture of the whole can be formed by imagining a pair of wrestlers. Each tries through physical force to compel the other to do his will; his *immediate* aim is to *throw* his opponent in order to make him incapable of further resistance. (1976, 75; emphasis in original)¹⁵

Thus every contest in disarmament leads to, and is motivated by, a subsequent contest in punishment, in which the winner of the contest in disarmament has an extreme bargaining advantage over the loser. Clausewitz claimed, however, that bargaining did not have to await the outcome of

13. Risk acceptance might also be caused by domestic political incentives—for an interesting discussion of this possibility in the context of World War I, see Goemans 2000. Evidence that aversion to losses makes people risk acceptant is emphasized by the experimental psychologists who developed prospect theory (Kahneman and Tversky 1979). However, the gamble associated with war involves not just a probable loss but a probability of a large gain combined with a larger probability of a large loss. The evidence that people are risk acceptant in those circumstances is not so clear. For a discussion of some of the pitfalls to avoid in thinking about attitudes toward risk, see O’Neill 2001.

14. In debates about the use of nuclear weapons during the cold war, a distinction was made between the countervalue and counterforce uses of weapons. Countervalue military contests are contests in punishment, and counterforce contests are contests in forcible disarmament. The first war in the Persian Gulf was an example of a counterforce contest. The ongoing conflict between Israel and the Palestinians is an example of a countervalue contest.

15. Here Clausewitz is talking about what he elsewhere calls “absolute war,” not “real war.”

the contest in disarmament but could precede it or accompany it. It should now be clear that, in the case of warring economic predators at any rate, there is good reason to take this claim seriously. To explain why *any* war occurs, therefore, one must explain why the adversaries could not have reached an agreement without fighting.

Bargaining and Fighting

It might appear that, if there is a range of divisions of disputed territory that two rulers both prefer to fighting over all of it, they will be able to agree on one of them rather than fight. However, when strikes occur it is obvious that there is a range of wage bargains that both labor and management prefer to shutting down the firm, or the industry, and yet strikes sometimes occur anyway. Thus while we have learned that there might often, perhaps even always, be compromise settlements that predatory rulers would prefer to fighting over disputed territory, that does not imply that they will in fact be able to agree on one without fighting.

In the case of strikes, as in any bargaining situation, the problem is that, while there are many agreements that both sides prefer to shutting down the firm, they have conflicting preferences about which of those agreements should be chosen, and strikes are a means of resolving that disagreement. But it is surprisingly difficult to explain exactly how a strike does that, why strikes are sometimes resorted to and sometimes not, or why some are so much longer and more costly than others. Explaining why wars occur is even more difficult.

As we saw in the previous chapter, attempts by economists to explain costly delays in reaching mutually beneficial agreements in situations resembling strikes have focused on the construction of explicit models of haggling (i.e., exchanges of offers and counteroffers that precede agreement). But there are several ways in which the bargaining process associated with war is more complicated than the one economists have focused on.

First, the essence of any bargaining process is the combination of a common interest in avoiding disagreement with conflicting interests as to the terms of an agreement. In the sort of bargaining situations exemplified by strikes, the disagreement outcome (a failure ever to agree) is fixed and is just an extension of the situation that exists while haggling occurs (in the case of strikes, the firm or industry is shut down). In the case of wars, the disagreement outcome (war) is not fixed but is the product of decisions made by the antagonists. Moreover, if Clausewitz is right, the disagreement outcome (absolute war) need not be the same as the war that is fought while the antagonists exchange offers and counteroffers (real war).

Second, the recent literature on bargaining by economists has focused on the role played by private information about the preferences of the bargainers, which they have a strategic incentive to misrepresent. In explaining wars, however, Clausewitz (implicitly) and Blainey (more explicitly) emphasize the role of conflicting beliefs about military capabilities, something that seems irrelevant to understanding strikes (though unregulated strikes, of course, have often been violent).

And third, the literature on bargaining typically assumes that the bargainers can be confident of getting any agreement they might accept. But the emphasis by structural Realists on the anarchic nature of international politics and the influential role the Prisoner's Dilemma game has played in shaping many people's beliefs about its implications make such an assumption problematic in any explanation of the occurrence of war. Indeed, some structural Realists would say that the fundamental cause of war is that agreements between or among states are unenforceable.¹⁶ In an intellectual environment that has been largely shaped by debates about structural Realism, one might almost say, paraphrasing the remarks by Clausewitz and Schelling about limited war quoted earlier, that were it not for Clausewitz it might be hard to get some people to take the subject of this chapter seriously.

If we are to follow the lead of the recent literature in economics about bargaining, we must deal with each of these complications. The obvious place to begin is to think about what Clausewitz called "absolute war," or a contest in disarmament. Even if Clausewitz is right in thinking that such wars rarely occur, this is the war that will take place if no agreement is possible, and therefore this is the war that the probabilities in figure 6 refer to. The nature of such a war will be determined by the strategies chosen by each side, but it seems safe to assume that each will choose what it believes to be the optimal strategy for disarming the other, given the expected strategy of its opponent, and therefore the properties of such a war can be assumed to be independent of the bargaining process.

Even so, there is no reason to think that all contests in disarmament are alike. But if we are to model the haggling process associated with absolute war, we must construct a model of haggling while fighting, and the nature of this process will be affected by how the war is fought. One of the disconcerting results of the economics literature on bargaining is that many conclusions about bargaining are dependent on seemingly minor properties of the process by which offers and counteroffers are exchanged. If that process is affected by how contests in disarmament are fought, we must be cautious about the generality of any conclusions we might reach that are based on assumptions about how a war is to be fought.

16. See, for example, Jervis 1978.

Because of the central role that Ariel Rubinstein's work on bargaining has played in the economics literature on the subject, a natural place to begin is to see if it could be applied to an analysis of bargaining that might take place during a contest in forcible disarmament.¹⁷ Rubinstein's model is based on two plausible assumptions: (1) bargainers alternate in making offers and counteroffers to each other, with the process ending when one bargainer accepts another's offer; and (2) they prefer agreements reached sooner to agreements reached later. One reason for the latter assumption might be that they discount future benefits, and another might be that delaying agreement would entail some risk that they would not be able to reach an agreement at all. Both seem potentially relevant to thinking about war.

Rubinstein showed that in a bargaining game that incorporates these assumptions there is only one Nash equilibrium that remains an equilibrium at every stage of the bargaining process.¹⁸ While a rigorous proof of this proposition is difficult, it is not so difficult to acquire an intuitive understanding of why it is true. Suppose the bargainers are negotiating over the division of a sum of money, and consider the possibility that an equal division might be an equilibrium. If this is an equilibrium at every stage of the bargaining process, then even if an equal division is not the opening offer, the other bargainer would counter with it and expect it to be accepted. But he could not get it until his turn came to make an offer, and therefore he should be willing to accept less than that now in order to avoid having to wait. Thus the assumption that an equal division is a subgame perfect equilibrium leads to a contradiction. To avoid such a contradiction, each bargainer must be indifferent between accepting what the other proposes immediately and getting his own demand one period later. There is only one division that satisfies this requirement, and it is the Rubinstein solution to the bargaining problem.

The importance of the requirement that an equilibrium continue to be an equilibrium at every stage of the bargaining process can be seen most clearly by thinking about prestrike negotiations. In such negotiations a bargainer who is dissatisfied with the most the other side is willing to offer can hope to do better only by shutting down the firm. But that would be costly for both sides, and at every stage thereafter either could make the other choose between accepting an offer or paying the cost of extending

17. For an exposition of Rubinstein's work, see his own account in Osborne and Rubinstein 1990. See also Muthoo 1999. For a nontechnical discussion of the economics literature on bargaining, see Muthoo 2000.

18. Game theorists call a Nash equilibrium with this property a "subgame perfect" equilibrium. As noted in the previous chapter, there are indefinitely many Nash equilibria in such a bargaining game, a fact that seemed for many years to imply that the concept of an equilibrium alone was not strong enough to imply anything about the bargaining problem.

the strike. The symmetry of their positions might suggest an equal division of the gains from agreement, but a bargainer who has the chance to make the first offer gains a slight advantage from the fact that the other would be willing to accept a bit less in order to avoid initiating a strike—an advantage that would alternate between them at every stage of the strike were one to take place.¹⁹

This reasoning implies that, if the amount of money to be divided and the extent to which each bargainer discounts future benefits are both commonly known, the bargainers should be able to reach an agreement without a strike: they already know everything they need to know to reach an agreement, and they also know that a strike would be costly but would not change anything. But if one believed that a strike would change what the other believed about one of these values, he might expect to get a more favorable agreement by striking. This, then, is a possible explanation of the fact that sometimes bargainers are able to reach agreement quickly and sometimes they are not. To complete this explanation, however, we would need to show how a strike could affect their beliefs.

Before considering that question, let us see whether this reasoning could be extended to a contest in forcible disarmament.

Fighting while Bargaining

Wars are unpleasant, which is reason enough to believe that warring rulers would prefer to reach agreements sooner rather than later. But unlike strikes, a contest in forcible disarmament can end before the combatants decide to end it: one or the other side might be defeated and therefore be unable to continue fighting. Part of the uncertainty associated with war is uncertainty about how long that will take, and therefore any delay in reaching agreement might entail not only the unpleasantness of further fighting but also some risk, however small, that the contest might reach a decisive conclusion before an agreement could be reached. A ruler who rejects an offer in hopes of getting a better one later, therefore, faces a risk that the war will be over before his demand can be accepted.²⁰ This is an additional reason for preferring to avoid postponing an agreement.

Of course, if the war ends there is a chance that a ruler might win, but there is also a chance that he might lose—it is that uncertainty that creates the possibility of an agreement in the first place. Postponing agreement is therefore a compound gamble: there is some chance that the war will end decisively before one's own demand can be accepted, and if it does there is

19. If the time between offers is small, this advantage will be small and the division will deviate only slightly from equality.

20. Think of the possibility that Saddam Hussein was holding out for a better deal at the onset of the second war in the Persian Gulf in 2003.

some chance that one might be defeated. Even so, exchanging offers while fighting occurs on the battlefield will often be feasible.

Thus there is reason to think that Rubinstein's analysis of bargaining is relevant to a contest in forcible disarmament (Clausewitz's "absolute war"). Note that Rubinstein's model implies that there is an advantage to being the bargainer who makes the first offer. In bargaining over the forcible redistribution of territory, it is obvious who that would be: the ruler who is dissatisfied with the current distribution would have to initiate a contest in forcible disarmament to change it, and therefore the satisfied ruler would have the advantage of making the first offer.²¹

But Rubinstein's analysis implies that, if both the information in figure 6 and each ruler's points of indifference between agreements now and agreements later are commonly known to both rulers, they should be able to reach an agreement without fighting. One has only to state that condition to see how difficult it is to satisfy it. But if Clausewitz and Blainey are right, a failure to satisfy it does not imply that they must fight an all-out contest in disarmament until one or the other is incapable of fighting further. Rather, war itself might reveal the information they need to reach an agreement, in which case war would be, as Clausewitz famously said, "simply a continuation of political intercourse, with the addition of other means." What we need to consider is how these "other means" could make an agreement possible if it were not possible at the outset.

One obvious possibility is that what happens on the battlefield reveals information about the combatants' relative military capabilities. In thinking about the significance of this fact, however, we must be careful to distinguish between two possible effects of battlefield outcomes on the expectations of the two rulers.

One is that they might become less uncertain about the ultimate outcome of the contest: as the contest progresses, it may seem more and more likely to both of them that one or the other will eventually win (as the outcome of a football game may seem less uncertain at the end of the third quarter than it did at the outset). But our discussion implies that this is irrelevant to the question of whether they could reach an agreement or not. As they become less uncertain of the outcome, the probabilities in figure 6 deviate more and more from equality. But as we saw, changing the probabilities has no effect on their interest in reaching agreement; it merely changes the terms of any agreement they might reach. Thus as the contest progresses one or the other might be willing to accept more and more unfavorable terms as it seems more and more likely that he will eventually lose, but the ability of the two rulers to reach an agreement

21. A model of absolute war with these properties is presented in Wagner 2000.

without fighting would not be affected by that fact, and anticipation of it prior to war is no more relevant than the fact that one might eventually win, since that information is already contained in the probabilities represented in figure 6.

What is important instead is that the rulers' expectations might become more *consistent*. If we label the two rulers i and j and call the probability with which each might win p_i and p_j , respectively, then consistency of their expectations requires that $p_j = 1 - p_i$, as is the case in figure 6. If this is not true, and the difference between p_i and p_j is great enough, then there may be no agreement they both prefer to fighting. One can readily see from figure 6, for example, that if they both expect to win with a probability of .8, then each would have to be given at least 60 percent of the territory if he were to choose not to fight, which is impossible.

But even if their expectations are not inconsistent enough to rule out the possibility of any mutually acceptable agreement, they may nonetheless be inconsistent enough to motivate fighting. For any inconsistency implies that each would expect experience on the battlefield to make the other less optimistic about winning and therefore willing to accept a less favorable agreement than he would be willing to accept prior to fighting. Thus while an agreement might have been possible prior to war, one ruler might expect to be able to get a better one by fighting, *while the other believed that to be unlikely*.²² This has a very important counterintuitive implication, which is that the possibility of ending a war with a negotiated settlement *might actually make war more likely than it otherwise would be*.

To see why this is true, look again at figure 6 and imagine that the expectations of the two rulers are somewhat inconsistent but not inconsistent enough to eliminate a range of possible divisions of the territory on the horizontal axis. Then if the status quo is within that range and it is not possible to reach an agreement after fighting begins, both rulers will prefer the status quo to fighting, and therefore neither would choose to fight. If, however, fighting does not rule out the possibility of subsequent agreement, then a ruler might expect that fighting for a while would reveal his true military strength and therefore lead to an agreement with more favorable terms.²³

On the other hand, if the rulers' initial expectations are so inconsistent that no agreement is possible prior to fighting, then the possibility of a negotiated settlement after fighting begins means that any war that occurs may be less costly than it otherwise would be, since it can be ended by mutual agreement if their expectations become consistent enough in the

22. If they both believed this to be true, their expectations at the outset would have been consistent, and therefore they would be able to reach an agreement without fighting.

23. Similarly, if strikes always led to the dissolution of the firm and never to agreements, one would expect fewer strikes to occur.

course of fighting. Thus the availability of negotiated settlements may make peace more or less likely, depending on the circumstances.

Before exploring the implications of this point, let us look more closely at exactly how battlefield outcomes might lead rulers to revise their expectations and what else, if anything, they might reveal.

Learning from Fighting

Learning from battlefield outcomes is based on inductive reasoning, an example of which was discussed at the very beginning of chapter 1: if one believes that it is more likely that an anonymous dog in the dog pound would resemble a Labrador retriever if it were a Labrador retriever than if it were not, then the fact that it resembles a Labrador retriever increases one's confidence in the hypothesis that it is one. If, on the other hand, it has characteristics that it would be expected to have if it were a pit bull, then one will be skeptical of the claim that it really is a Labrador retriever.

Similarly, it is plausible to think that military leaders begin a war with an idea of how it will be fought, which leads them to think that certain battlefield outcomes are more likely than others. If these expectations are borne out, then their confidence in them will increase, but if not they will decrease.

As noted in chapter 1, this reasoning can be justified by the axioms of probability theory. One's expectations prior to receiving new information are represented in a "prior" probability distribution (like, e.g., the ones in figure 6), and new information leads to a revised, or "posterior," probability estimate, which is a conditional probability: the probability that one's hypothesis is correct, given that the event in question occurred. For this to be possible, two other conditional probability estimates are required: the probability that the new event would have occurred if one's hypothesis was correct and the probability that it might have occurred if one's hypothesis was incorrect. The formula that allows one to compute a posterior probability from this information is called Bayes's rule, and the process is called Bayesian updating. Bayes's rule implies the relation between prior and posterior probabilities just described.²⁴

Note that all these probabilities are based on some understanding of how the war will unfold and not just on knowledge of the number of some objects in a larger universe (like the number of aces in a deck of cards). But, like scientific theories, this understanding will be based on both creative guessing and deductive reasoning. Thus learning can consist not merely of Bayesian updating but also of the discovery of possibilities that

24. An account of all this in the context of scientific reasoning can be found in Howson and Urbach 1993. People often make mistakes in such reasoning, which is an important example of how human decisions may not be a "reflective equilibrium."

one had not thought of. For example, in the second war in the Persian Gulf the advance of U.S. forces in Iraq apparently was initially more difficult than expected because the Iraqis decided to use irregular forces to attack the extended supply lines of U.S. troops as they advanced toward Baghdad. But newspaper accounts indicate that this was an Iraqi strategy that U.S. military planners had not anticipated. It is important also to note that there were disagreements among military commentators about how much revision this unexpected development required in initial U.S. expectations about the eventual outcome of the war.

In discussing how Rubinstein's bargaining model might be extended to an analysis of bargaining while fighting, I implicitly assumed that the military contest proceeded continuously in the background while rulers exchanged offers and counteroffers and that victory or defeat could come at any time. However, while that might be true of the final stages of a military contest, this discussion of learning from battlefield outcomes calls attention to the fact that many wars are made up of discrete battles, and battles early in the contest do not entail much risk of total defeat. Such battles can have two effects: (1) they can change the probability that one side or the other will eventually win, and (2) as just noted, they can convey information about what those probabilities are.

But if a battle is fought with the second effect in mind, then bargaining will be delayed until its outcome has been observed. Moreover, this would continue to be true until both sides thought that no more favorable information could be conveyed by further fighting. This perhaps helps explain why, even though there may have been prewar attempts to reach a negotiated settlement that failed, once a war begins peace negotiations typically do not occur throughout military conflicts but are resumed only toward their end. As Paul Pillar wrote in an important study of peace negotiations, "the opening of peace negotiations usually must await a common perception of the trend of military events" (1983, 199). Thus while an exchange of offers and counteroffers could occur throughout a military contest, it usually does not.

Moreover, while battles can both change the balance of military capabilities and convey information about it, it is possible for battles to be fought whose only function is to convey information. Such battles can occur even in the midst of a contest in disarmament. Consider, for example, General von Falkenhayn's discussion of the German military position at the end of 1915, during World War I, which includes the following passage describing plans for the battle of Verdun:

the strain on France has almost reached the breaking point. . . . If we succeeded in opening the eyes of her people to the fact that in a military sense they have nothing more to hope for, that breaking point would be reached. . . . To achieve that object the uncer-

tain method of a mass break-through, in any case beyond our means, is unnecessary. We can probably do enough for our purposes with limited resources. (Falkenhayn 1920, 249)²⁵

But it is clearly also possible that battles might convey information about relative military capabilities even though they are not part of a military contest that would eventually lead to the complete disarmament of one side or the other. This appears to be what Clausewitz had in mind when he wrote that in “real wars” states might not try seriously to disarm each other at all.

The possibility of revealing information without actually attempting to defeat one’s adversary implies, as Clausewitz suggested, that the contest I originally described might not occur at all, even though expectations about its outcome would motivate any agreement that might be reached. Thus the wars that we see are not necessarily good guides for constructing a model of a war that would be fought if the aim were only the complete disarmament of the enemy. And even if such a contest began, peace negotiations might take place in the context of a cease-fire agreement, which could be accepted because both sides thought there was no further information to be revealed by fighting. This reinforces my earlier comment about the difficulty of constructing a truly general model of bargaining and war.²⁶

Revealing Private Information

The preceding discussion of bargaining while fighting is very different from the literature about bargaining that has been developed by economists. To understand why, we must look again at figure 6. There are two elements of figure 6 that will influence the terms of an agreement that the two rulers might accept as an alternative to war. One is their probabilities of winning the contest, and the other is the shapes of their indifference curves. We have just seen that if they have inconsistent estimates of their probabilities of winning and believe that fighting could change them, they could have a motive to fight for a while in order to improve the terms of a deal. We must now consider whether the same might be true of their indifference curves.

Part of the answer to that question is the same as the one just given about the probability of winning: battles also convey information about the costs that a contest in disarmament would entail.²⁷ As already noted,

25. I owe this example to Hein Goemans.

26. There are now a number of different models of bargaining while fighting. For summaries of recent work on the subject, see Powell 2002 and Reiter 2003.

27. This perhaps helps explain why the point of military operations might be, as Clausewitz, said, “neither to conquer the enemy country nor to destroy its army, but simply *to cause general damage.*”

while the costs of fighting are not explicitly represented in figure 6, they will influence the shapes of the indifference curves that are portrayed there. And therefore battlefield outcomes can influence not only the probabilities on the vertical axes but also the points of indifference between probabilities and territorial divisions.

However, two rulers with the same expectations of the consequences of fighting a contest in disarmament, including expectations about its costs, might nonetheless have different points of indifference between territorial divisions and probabilities of winning. Moreover, by misrepresenting his preferences a ruler could improve the terms of the agreement. The question we must consider, then, is whether fighting would also be a way of revealing information about the true preferences of the two bargainers. That is the question that has been the main focus of the literature on bargaining developed by economists.²⁸

A technique for constructing game theoretic models of this problem was developed by John Harsanyi, for which he received a Nobel prize in economics (Harsanyi 1967–68). It is based on the idea that, while a player of a game might not be certain about the preferences of another player, he might nonetheless have an idea of what range of possible preferences he might have and be able to assign each of them a probability. If the second player knows that the first has only probable knowledge of his preferences, then whatever the second player's preferences actually are, his choices would be influenced by his knowledge of the first player's uncertainty. Thus one might imagine that the uncertain player confronts one of many possible players, each of whose choices would be determined by his true preferences, and those choices might therefore reveal information about what that player's preferences actually are. Each of the possible sets of preferences that a player might have is commonly said to determine his "type."

Clearly, learning from the choices made by someone who has an incentive to mislead you is a far more complex problem than learning about an adversary's military capabilities by observing battlefield outcomes. Battlefield outcomes cannot be faked, and therefore the probability with which they were expected to occur is only a function of one's prior understanding of how the war would be fought, how weapons would work, how well trained and motivated the military forces were, and so forth. The probability with which an adversary would be expected to make

28. For a survey, see Kennan and Wilson 1993. Note, however, that in bargaining in the context of war the preferences of more people than the ones conducting the bargaining are relevant. Information about the reactions of domestic political actors, or the decisions of potential allies, is not revealed by learning about the preferences of the people conducting the bargaining but can be garnered by observing the course of the war just like information about military capabilities.

a choice one can observe, however, will be a function both of whatever his true preferences happen to be and the effect he expects his choices will have on one's own expectations. If he is also known to be uncertain about one's own true preferences, then the problem is even more complex and there are often multiple equilibrium combinations of choices.

Many game theorists would claim that I have exaggerated the difference between learning about preferences and learning about relative military capabilities. The reason for this claim is that, if two people with common prior probabilities always update them in accordance with Bayes's rule, then their probabilities will remain the same. Any differences in their probabilities must therefore be the result of their having been exposed to different information, and therefore if I learn that another person's probability estimate is different from my own I should infer that he knows something I do not know and change my estimate accordingly. Thus, in the situation represented in figure 6, both inconsistent probability estimates and inconsistent beliefs about each other's true indifference curves would be the result of private information that the two rulers had an incentive to conceal, and both could therefore be revealed by the choices they made during any bargaining that accompanied fighting.²⁹

Clearly one's enemy may have an incentive both to feign confidence in his military capabilities that he does not really have and to conceal some of them so that one cannot take countermeasures against them. And therefore knowledge that he is unexpectedly confident of winning a military contest ought to make one wonder if he knows something one does not know. That does not imply, however, that his bargaining behavior is a perfect substitute for battlefield outcomes as a source of information as to his true military capabilities, since even if one had access to all the information he had one might disagree with the inferences he drew from it. This is not only because one might doubt his competence as a military strategist but also because, like science, learning about war is only partly a matter of Bayesian updating. A military contest is like a very costly experiment that tests competing theories about how the war will unfold, and just as scientists with different theoretical commitments frequently disagree about what they think experiments will reveal, so equally well-informed military strategists may disagree about what battlefield outcomes to expect.³⁰

29. See, for example, the discussion of this question in Fearon 1995b.

30. This is an example of an issue one would be unlikely to think of if one had not subjected one's thinking to the discipline of a formal model. For a discussion of the relevant theory, see Geanakoplos 1989 and 1992. For an interesting discussion of war as a way of testing competing theories, see Smith and Stam 2004. For accounts of competing bets and heated disputes between scientists backing different hypotheses, see Glanz 1998 and 1999. See also the discussion from a Bayesian perspective of the strength of scientists' beliefs in the truth of their ideas in Press and Tanur 2001.

A plausible reading of Pillar's book on peace negotiations is that statesmen rely primarily on battlefield outcomes for information about military capabilities and then reveal any remaining private information during the bargaining process that accompanies peace negotiations (Pillar 1983). However, while Pillar interprets war termination as a bargaining process, his book was written before models of bargaining with incomplete information were available and therefore does not actually investigate this question. There are many historical studies that criticize statesmen's failure to learn rapidly enough from battlefield outcomes (e.g., Iklé 1991) but few careful studies of exactly how they do it. This is an important research frontier in the study of war.³¹

The theory of games with incomplete information had a big impact on the literature about deterrence during the cold war, where the issue was how political leaders could reveal prior to war whether they were really willing to carry out deterrent threats if they were challenged. Much of that literature assumed that bargaining ended when war began and therefore whatever war had been threatened would occur if a defender who was not bluffing was challenged.³² This overlooks the fact that, as Clausewitz said, "war in itself does not suspend political intercourse or change it into something entirely different." However, the cold war literature on limited war is consistent with Clausewitz's idea that the function of limited war is to reveal information about absolute war, which might mean that absolute war never occurs.³³

Bargaining, War, and Alliances

Like most discussions of both war and bargaining, the analysis so far has been couched in terms of a contest between only two predatory rulers. But wars can involve more than two states. We must therefore consider what effect adding more rulers would have.

With more rulers, alliances become possible. Alliances can affect not only the conduct of war but also the bargaining process that might accompany it. In thinking about these added complications, I will follow the same analytical strategy employed earlier and consider first what effect they would have on a contest in disarmament (Clausewitz's absolute war)

31. For a pioneering study of how German and French leaders responded to the same course of events on the battlefield during World War I, see Goemans 2000.

32. This assumption is reflected in Fearon's (1995b) pioneering article on this subject.

33. For example, one of the main themes of the recent history of the Korean War by William Stueck (1995) is that the Korean War can usefully be considered to have been a substitute for World War III.

and then introduce the possibility of bargaining over the terms of a negotiated settlement that might be accepted as an alternative to such a contest.

The idea of subjective or personal probabilities seemed sufficient as a way of capturing the uncertainty associated with a contest in disarmament between two rulers, since they contained all the information that was important in making a choice between fighting such a contest and accepting a division of the territory in dispute instead. However, in thinking about contests among varying configurations of allies it will be necessary to think about the effect of shifts in alliances on the probability with which one side or the other would be expected to win. Shifts in alliances would lead to realignments of the military forces that would fight each other, and it is the distribution of those military forces that would determine the degree of confidence that a ruler would have about his ability to disarm his adversary. So if we are to think about the effect of alliances we must say something about the effect of any particular distribution of military forces on the probability of winning or losing.

As I pointed out earlier, personal probabilities are both subjective and nonarbitrary: they reflect the choice that a person would make between the uncertain prospect he actually confronts and a lottery with the same outcomes and known probabilities, but these choices would obviously be based on everything that person knew that he thought might affect which outcome occurred. Thus there are two potential sources of disagreement about the effect of any particular distribution of capabilities on the expected outcome of war: there might be disagreement about what capabilities were relevant or how they should be measured and disagreement about what any particular distribution of capabilities implied about the probability of winning or losing.

While any assumptions we make about these issues will be arbitrary, we must make some assumptions if we are to think about the relation between alliances and war. I will therefore make assumptions that reflect the way these issues are often discussed, while bearing in mind that they are arbitrary. If such assumptions lead to conclusions that differ from claims commonly made by writers on the subject, they can provide the basis for a counterexample to those claims. But before leaping to any conclusion as to what the right answer to the question really is, we would need to consider whether different assumptions would lead to different conclusions.

To get the analysis started, then, I will assume (1) that military capabilities can be measured at least to the extent that one can determine the ratios between them (so that one can say, for example, that one side has twice the military capabilities of the other) and (2) that the ratio between the probabilities with which each side might win a military contest is the same as the ratio between their military capabilities (so that, e.g., if one

side is twice as strong as the other, it is twice as likely to win a contest in disarmament between them).³⁴

Thus, if we label the military resources of state i as r_i and the probability that one state will disarm another p , in a two-state contest

$$\frac{p_i}{1 - p_i} = \frac{r_i}{r_j}.$$

This implies that

$$p_i = \frac{r_i}{r_i + r_j}.$$

And therefore the probability of victory of each state can be equated with the percentage of total military resources that it controls.

To make things as easy as possible, let us assume there are only three predatory rulers who might participate in a contest in disarmament. If one attacks another, the third could either join in or not. If it were to join the fight we must consider how that would affect the probabilities associated with the outcomes. One possibility is that, if two rulers both fight the third, then the third state faces the sum of the military resources of the other two, and therefore the probability that this lone state k will win will be

$$\frac{r_k}{r_i + r_j + r_k}.$$

The probability that the other two will be victorious will, of course, be the complementary probability, which implies that, if two of three equally powerful states fight together, they will be twice as likely to defeat the third as each would be separately.

But what happens if the two rulers who fought together succeed in disarming the third? Many writers assume, implicitly or explicitly, that they will then divide the territory of the defeated ruler between them. But if they could do that, one might wonder why there could not have been an agreement dividing the disputed territory among all three rulers at the outset. Much of the literature assumes that this is not possible.³⁵ But this begs the question raised by Clausewitz's analysis of war, which is why states cannot reach negotiated settlements without fighting. It makes more sense to ask

34. For references to the literature about contests that conform to these assumptions, see Skaperdas 1998.

35. See, for example, the important recent contributions to the literature by Schweller (1998) and Powell (1999, chap. 5), both of which explicitly make this assumption.

first how a contest in disarmament would be fought and then to consider what negotiated settlements might be feasible as an alternative to it.

If an absolute war is a contest in disarmament fought until one state has rendered its adversary incapable of further resistance, and the point of such a contest is to enable a predatory ruler to capture all the territory in dispute, then if there are three rulers an absolute war would be a contest in disarmament fought until every state but one had been disarmed and one ruler therefore controlled all the territory. If two rulers fought a third and disarmed him, therefore, then the contest would not be over until they had fought each other.

The probability that state i will eventually disarm the others if it begins as state j 's ally is the probability that states i and j will win the first stage of the contest times the probability that state i will defeat state j in the second stage. In the following analysis I will initially assume that defeat entails the destruction of the third state's military capabilities—for reasons that will become clear, the possibility that the defeated state has military resources that can be transferred to the victors after its armed forces have been destroyed will be examined separately in the next chapter. If defeat entails the destruction of a state's military capabilities, then the relative power of the two victorious allies will be unchanged by the defeat of the third state.³⁶ The probability that state i will eventually disarm the others if it begins as state j 's ally is therefore

$$\frac{r_i + r_j}{r_i + r_j + r_k} \left(\frac{r_i}{r_i + r_j} \right),$$

which reduces to

$$\frac{r_i}{r_i + r_j + r_k}.$$

Thus with these assumptions each ruler faces the same probability of defeating the other two whether it fights alone or fights with an ally in the first round of the contest. Each would therefore be better off sitting out the first round and then challenging the winner, since no matter what the distribution of military resources, the probability with which any state will win a contest against either of the other two will be greater than the probability with which it would win a contest against the other two combined.

36. It is possible that fighting the first round might weaken the allied states, but so long as neither expects to be affected proportionately more than the other this would not affect their prewar expectations.

But the other two would then prefer to join together to fight such a state rather than fight each other first, and that is therefore how we should expect this contest to be fought.³⁷

It is instructive to compare this conclusion with Kenneth Waltz's famous claim that the potential danger posed by a powerful ally would lead weak states to ally with each other to "balance" the power of stronger ones rather than to "bandwagon" with stronger states against weaker ones. As we saw in our discussion of Waltz's claim in chapter 1, the argument Waltz gave for it is incomplete. A state forced to choose between a stronger and a weaker ally confronts a complex tradeoff: if it joins the stronger side it will confront a more powerful adversary after victory, but if it joins the weaker side victory will be less likely.³⁸

In the little model of absolute war just discussed, these two factors exactly cancel each other out, and therefore states should be indifferent between the two possibilities. Moreover, if powerful allies can be expected to bear a larger proportion of the costs of war than weaker ones, then a state would actually prefer a more powerful ally.³⁹ This example illustrates once again how easy it is to overlook the implications of even simple assumptions.

Nonetheless, the implications of these assumptions seem very counterintuitive. Before accepting the implications at face value, we need to consider whether there is something wrong with the assumptions.

The conclusions we arrived at are the result of the way military resources are assumed to accumulate as compared to probabilities: military resources are added together to produce greater strength, whereas probabilities are multiplied. That is why, if all three rulers have equal resources, a ruler who fights both the other two together faces the sum of their resources and therefore has a one-third probability of winning, whereas if he fights them separately he has a fifty percent chance of winning each contest, but the accumulation of risk implies that he has a probability of beating both of them separately of only .25.⁴⁰

If alliances combined military capabilities in a way that was not sim-

37. If the states all have the same military capabilities, for example, then in a contest in which two first fought the third each would have a probability of winning of one-third, whereas a state that sat out the first round would have a probability of defeating the survivor of .5. But then each of the other two would have a probability of defeating both the others of only .25, and they could therefore do better by joining together to fight the third before fighting each other.

38. For Waltz's argument, see the discussion in chapter 1. See also the discussion of it in Powell 1999, chap. 5.

39. This is one of the assumptions in Powell's analysis of this question (1999, 162).

40. Moreover, a state that combines its military forces with its allies' forces increases the probability of winning the first round of the contest but exposes itself to the risk that it might be defeated along with its ally.

ply additive, that is, if there were economies of scale in alliances, fighting an alliance would be more risky. Using the notation in Powell 1999, let g be a parameter that describes the possible effect of economies of scale in alliances. Then the probability that state i will win a military contest if it initially allies with state j will be

$$\frac{r_i}{r_i + r_j} \left[\frac{g(r_i + r_j)}{g(r_i + r_j) + r_k} \right],$$

which reduces to

$$\frac{gr_i}{g(r_i + r_j) + r_k}.$$

If $g > 1$ and $r_j > r_k$, then the value of this expression would be increased if state i allied with state k instead of state j , and now Waltz's claim would be supported.⁴¹

However, even if there are economies of scale in alliances, they might not be great enough to outweigh the advantage of waiting out the first round of the contest and just fighting the winner. Moreover, it is also possible for coalition warfare to be inefficient, so there might be diseconomies of scale in an alliance. If so, a weak state would actually prefer to fight alone.⁴² Thus there are many possible combinations of factors that might, in any given situation, lead a state to prefer balancing to bandwagoning, to have the opposite preference, or to be indifferent between them, and therefore, contrary to Waltz, no general statement can be made as to what states will do.

In talking about absolute war in a two-state setting, I pointed out that, while the optimum way of conducting such a war posed a complicated strategic problem in its own right (since each state's strategy depended on the expected strategy of the other side), there was good reason to think that the solution to that problem was independent of any bargaining that might take place between the two adversaries. Expectations about the outcome of that contest could then be taken as the disagreement outcome in bargaining over the terms of a negotiated settlement.

This discussion of a three-way contest in disarmament implies that, as one might expect, expectations about its outcome depend not just on the military capabilities of the adversaries and the strategies they employ but also on information about the effect of alliances and perhaps also about which alliances will form. Nonetheless, the solution to the problem of how

41. On this point see also Skaperdas 1998.

42. For a possible example, see the statements of British policymakers about defending France in the period before World War II quoted in Schweller 1998, 150.

to fight a three-way contest in disarmament is also independent of any bargaining that might occur, and, given the other requisite information, expectations about its outcome can also be represented as a set of probabilities that express how optimistic each state would be about eventually emerging as the sole winner of such a contest. Moreover, if all states are equally matched they would each be far less optimistic about winning such a contest than if there were only two states.

In the simple case we first looked at, for example, if all states have the same military capabilities they would each expect to win a contest in disarmament among them with a probability of one-third, and therefore each might be willing to accept less than one-third of the territory in dispute rather than fight for all of it. This is because, while two of them could, with a much higher probability, capture all the territory, they would still have to fight each other for control over it. Of course, they might anticipate being able peacefully to agree about how to divide it, but if so there seems to be no reason all three states could not do the same. That is the question we must now examine.

Unfortunately, the analysis of n -person bargaining is much less developed than the analysis of two-person bargaining. Moreover, it is even more difficult to construct a model of bargaining while fighting when there are more than two states than when there are just two. Perhaps the best we can do is construct a plausible scenario and identify this as an important question for future research.

Suppose, then, that a contest in disarmament among three states begins with a contest between two of them against the third. If the two allied states win, then the second stage of this contest will be a two-state contest like the one examined previously. With complete information the two victorious states will therefore agree to a division of the territory without fighting.

Now suppose that bargaining during the first stage of the contest takes place as follows: one ruler proposes a division of territory to the other two. Each in turn can accept the offer or propose another. If both accept then the conflict ends and the division is implemented. If a ruler whose turn it is to respond proposes another division then the contest continues, and if no one has been defeated by the next period the other two respond to that proposal. The contest continues in this way until one side has been disarmed or all the participants accept a division of the territory.

This is a scenario that resembles the Rubinstein two-person bargaining game. There is a three-person version of Rubinstein's game that has been discussed in the literature, in which there is a subgame perfect equilibrium set of offers similar to the ones that characterize the two-person bargaining game. As in the two-person case, therefore, with complete information one might expect that agreement would be immediate and

therefore the military contest would not occur.⁴³ However, if expectations about the contest were not consistent, then there might be an incentive to fight more limited wars to reveal information about what to expect should an absolute war be fought. Thus the relation between bargaining and war in a world of three states appears to be qualitatively similar to the relation when there are just two.

There is one striking difference between a three-state world and the two-state case discussed earlier, however, which is that in the three-state case a war of all against all is the disagreement outcome in any bargaining that may occur even if the only wars that occur are bilateral wars. This is because a war between state *i* and state *j* that reveals information about state *j*'s military capabilities has implications for the terms of a negotiated settlement involving a possible war among all three states. Thus the outcome of a war between states *i* and *j* could lead to a revision of the territory held by state *k*, even though state *k* did not participate in the conflict. In this situation everything concerns everybody, whether they all participate in a military conflict or not.

Bargaining, War, and the Balance of Power

We saw in chapter 2 that in Western political thought the concept of a state system dates back at least to fifteenth-century Italy and that what Kenneth Waltz called "balance-of-power theory" has long been an important element of thinking about it.⁴⁴ One controversy about state systems concerns how to explain the ability of states to maintain their independence. Another concerns the effect of systems of independent states on human welfare. The most important issue in the latter context is the frequency of warfare, though the Prisoner's Dilemma game has led many people to doubt the ability of independent states to cooperate in the pursuit of any common interest. What Waltz called balance-of-power theory includes controversies about how to explain both the ability of states to maintain their independence (often called the stability of state systems) and variations in the frequency of warfare among them.

Sometimes in these controversies the balance of power refers to the

43. For a discussion of this game, see Osborne and Rubinstein 1990, 63–65. In the three-person game, any convention about how goods should be divided can be supported by strategies that reward someone who rejects a deviant offer by giving him all the gains from bargaining in the following period, and therefore subgame perfection does not guarantee uniqueness in the three-person game. But it is not clear in this context where such a convention might come from. However, for the purposes of this discussion it does not matter what the outcome of this bargaining game is expected to be, so long as, with complete information, there is an equilibrium outcome.

44. For general surveys of writings on this subject, see Claude 1962 and Sheehan 1996.

distribution of military capabilities among individual states, and sometimes it refers to the distribution of capabilities between alliances. Common sense suggests (wrongly) that war between two states is least likely when their military capabilities are equal. Often it is assumed implicitly that if states are to maintain their independence, then weak states must “balance” against the power of strong ones. Common sense might also suggest that balancing would serve to reduce the likelihood of war when there are more than two states, though many writers have denied this.

It seems likely that the availability of negotiated settlements as an alternative to war will have an impact on both the frequency of war and system stability. We have seen that negotiated settlements that everyone prefers to war may not always exist, and in the next chapter we will see that the necessity that agreements be self-enforcing may reduce further the number of agreements that are feasible. Nonetheless, it is obvious that states often do accept such agreements and that Clausewitz was right in claiming that the possibility that they might be accepted has a profound impact on both the likelihood of war and how it is conducted. Before looking at factors that may restrict the availability of negotiated settlements, therefore, let us consider what effect their availability might have on controversies about the balance of power and its significance.

The Distribution of Power and the Likelihood of War

Common sense suggests that war between two states is least likely when their military capabilities are equal, but commonsense reasoning fails to take into account the effect of fighting on the bargaining that accompanies it. Donald Wittman (1979) argued that, since the distribution of military capabilities would affect the terms of any agreement they might accept but not whether they both would prefer an agreement to war, the balance of power should have no effect on the likelihood of war. Geoffrey Blainey (1988) claimed, however, that states would be more likely to agree on the terms of an agreement if their capabilities were unequal than if they were equal. As we have just seen, the fact that a mutually acceptable agreement exists does not mean that states will not fight, so Blainey might be right. However, the argument that he gave for his claim was both incomplete and confused, so we must decide for ourselves whether there is any reason to believe that it is true.

Blainey’s fundamental claim was that “[w]ars usually begin when two nations disagree on their relative strength, and wars usually cease when the fighting nations agree on their relative strength” (1988, 293). The reason he gave for this claim was that disagreements about relative strength lead to disagreements about relative bargaining power (115–19). As we have seen, there is good reason to take this claim seriously.

The reason Blainey thought that war was more likely when power was distributed equally than when it was not was that he thought states were most likely to disagree about their relative strength when they were equal (108–24). One has only to state this claim to see a serious problem with it: if states disagree about their relative strength, how can one say whether they are equal or not?

The only reason Blainey gave for believing this proposition was true was that warring states find it easier to reach agreement after fighting than before fighting began—otherwise the war would not have occurred. But when a war ends, Blainey claimed, one state has revealed that it is stronger than the other. From this he concluded that peace was most likely when there was a “clear preponderance of power” (113).

Not only is this not a valid argument, but its plausibility rests on a confusion between *uncertainty* and *inconsistent expectations*. As we saw, as a war progresses states may become less uncertain about how it will end, but this implies nothing about whether the two states’ expectations are consistent—indeed, if they are both maximally uncertain their expectations *must* be consistent, since they would both assign equal probabilities to winning and losing.⁴⁵ Since it is the consistency of their expectations that is important in reaching agreement, states that are maximally uncertain about the outcome of war may have no difficulty in agreeing on the terms of a negotiated settlement they would both prefer to fighting, if a negotiated settlement is possible at all. This would be compatible with another intuition that is common in the literature, which is that negotiated settlements are most likely in wars that are stalemated.⁴⁶

Nonetheless, it is possible to construct an argument in support of the proposition that equality of power can make war more likely. But doing so reveals that the proposition is not always true.

As we have seen, the balance of power between or among states can be thought of in two ways: as a distribution of subjective probabilities assigned to the possible outcomes of a contest in disarmament between or among those states and as a measure of the distribution of the military capabilities among them on which such probability estimates would be based (e.g., size of armed forces, nature and number of weapons systems, size of population, quantity of industrial production, and so forth). How-

45. However, the consistency of their expectations might not be common knowledge, so it would still be possible for one ruler to feign confidence he did not really have.

46. I suggested earlier that a stalemate is best thought of in terms of expectations about the length of a military contest rather than which side is more likely ultimately to win. Nonetheless, one reason a war might be expected to last a long time is that the two sides are thought to be evenly matched. The importance of a “mutually hurting stalemate” in producing negotiated settlements of civil wars has been emphasized by William Zartman. See his essay in Licklider 1993 and the discussion of that thesis in other essays in that volume.

ever, some of the factors that might be expected to influence the outcome of a contest in disarmament are more easily identified and observed than others. In addition to the factors just mentioned, for example, the training, morale, and fighting spirit of a state's armed forces, as well as the nature of the strategies that will be employed in fighting, are also important. Thus some of the factors that determine the probability with which a state will win a military contest are more easily observed than others.

It is plausible to think that the more evenly matched two states are with respect to the military capabilities that can be observed and measured, the greater the significance of the factors that cannot be so easily observed. Just as the quality of the coaching may determine the outcome of a professional football game but have little effect on the outcome of a football game between a professional and a high school team, so the outcome of a war between two equally powerful states may be determined by which side has the better generals, but not even the best general could enable Guatemala to disarm the United States. And it is easy to see how each of two rulers could believe in his own military genius even though it was common knowledge that their military forces were evenly matched. This provides a possible justification for Blainey's claim and also solves the problem of how one could say that two states were evenly matched even though they disagreed about their relative military capabilities.

However, we have seen that, if rulers are to reach a negotiated settlement, they must have consistent beliefs not only about their probabilities of winning a military contest but also about the cost of doing so. And it is clearly possible for rulers who agree that one is much more likely to be able to disarm the other nonetheless to disagree about the cost the stronger state would have to pay. The U.S. war in Vietnam may be an example of this: it was perhaps the great disparity in military power between the United States and North Vietnam that made political leaders in the United States underestimate the ability of North Vietnam to impose costs on the United States, and the ability of North Vietnam to force a revision of that estimate led to a settlement of the contest that was far less favorable to the United States than the one it had expected at the outset.⁴⁷

Of course, inconsistent expectations about a contest in disarmament are not sufficient for fighting to occur. A ruler must also be optimistic enough about his ability to alter the expectations of his opponent at an acceptable cost to make it worth his while to try. Thus "real war," like "absolute war," is a costly contest with an uncertain outcome. But the probabilities assigned to the possible outcomes of real war need have little relation to the probabilities assigned to the outcomes of absolute war.

47. The bargaining process that accompanied the termination of the war in Vietnam is examined at length by Paul Pillar (1983) in the book referred to previously. Pillar interprets the war in Vietnam as primarily a contest in the imposition of costs.

Even if a ruler hopes to alter the enemy's beliefs about his own relative military strength, the aim may not be to show that he is stronger than the enemy but merely to show that he is not as weak as the enemy thought. And it is possible that the point of the contest will not be to affect the enemy's estimate of the probability of winning the contest at all but to influence his estimate of the costs that winning would entail. Thus it is possible to hope to gain a bargaining advantage by losing battles, which helps explain why Clausewitz claimed that weak states could hope to gain by fighting stronger ones and why military forces might be used not to disarm one's adversary but merely "to cause general damage."

We saw that what is true of two-state contests seems likely also to be true of multistate contests. Thus Wittman was right in claiming that there is no general connection between the distribution of military capabilities between or among states and the likelihood of war.⁴⁸

The Distribution of Power and the Independence of States

As already noted, the phrase *balance of power* is sometimes used to refer to the distribution of military capabilities among individual states and sometimes to the distribution between warring alliances. When it refers to individual states, it often has meant not that the individual states were equally powerful but that no individual state was powerful enough to defeat all the others combined. As we saw, that was how Saint-Pierre used the term.⁴⁹ Similarly, Friedrich Gentz wrote early in the nineteenth century that "if the states system of Europe is to exist and be maintained by common exertions, no one of its members must ever become so powerful as to be able to coerce all the rest put together" (quoted in Gulick 1955, 34).

If this condition is satisfied, weak states could perhaps preserve their independence by joining together to oppose strong ones. And if they do, then the power of strong states will be balanced by the power of an opposing coalition. This is how the distribution of power among individual states, "balancing" (to use Waltz's influential terminology), and the distribution of power between opposing coalitions are related.

As the seventeenth-century tract by the Duke de Rohan discussed in chapter 2 illustrates, the history of Europe can plausibly be told as a history of states forming balancing coalitions to oppose attempts by power-

48. Note that this discussion has been based on the assumption discussed previously that states rely primarily on battlefield outcomes in revising their expectations about their relative military capabilities. The question of the relation between the distribution of power and the likelihood of war is much more complex if one assumes that each side knows its own true capabilities but misrepresents them and that that information is revealed in the course of making offers and counteroffers.

49. See the discussion in chapter 2.

ful states to establish hegemony over them: first Spain, then France under Louis XIV and Napoleon, then Germany under the kaiser and Hitler, and then the USSR after World War II. Moreover, the tendency for such balances to form could be offered as an explanation of the fact that none of those attempts was successful.

However, those would-be hegemonies had allies. Moreover, while no European state succeeded in eliminating all the others, balancing did not protect the independence of states in the ancient Chinese Warring States System, the Greek city-state system, or the subsequent Hellenistic one.⁵⁰ And, in spite of what Waltz claimed about balance-of-power theory, no one has offered a valid argument for the proposition that weak states should always be expected to ally with each other against strong ones.

In addition, it is not even clear why balancing would protect the independence of states. It might do so if the equality of power between coalitions prevented war from occurring, since if no war occurred no state could be defeated. But the frequency of war in the European state system seems to rule out that possibility. Moreover, we have already seen that the idea that equality of power between antagonists makes war less likely than it otherwise would be is itself based on invalid reasoning. But if wars between evenly matched antagonists occur, then one would expect that at least some of the time the potential hegemon would win. And even if it lost, one must ask why it would not be eliminated by the members of the victorious coalition.⁵¹

In Europe, balance of power thinking can be traced back at least as far as the Renaissance, when warring princes competed for control of northern Italy, and it flourished in the eighteenth century. That was a time when, as Rousseau's writings illustrate, international politics could plausibly be said to have consisted of struggles among predatory rulers for control over valuable territory. In those circumstances, as we have seen, everyone is in conflict with everyone else, even if they might ally temporarily, but compromises are possible because territory is divisible. This implies a different relation between the balance of power and the ability of states to maintain their independence.

There are two ways in which states might lose their independence: they might be disarmed in a military contest, or they might agree to give up their independence in a negotiated settlement. But these are in reality not two ways but one, since disarming a state only weakens its bargaining power but does not determine what will happen to it. Thus whether any

50. For a comparison of the history of Europe with the history of China that focuses on this question, see Hui 1999 and 2000.

51. This last question is perhaps why some authors have claimed that balancing is something that only satisfied states would engage in. However, as we saw, Waltz explicitly denied that this was true.

particular distribution of military capabilities among states leads to a loss of independence for one of them depends on both the preferences of their leaders and their bargaining power.

If rulers are warriors trying to maximize their ability to profit from the labor of others and there are economies of scale in predation, then they might all be better off if they joined forces, and therefore any conflict between them would only be about the terms on which they would give up their independence. All three states might therefore disappear even though they were all militarily equal. Thus in Europe many of the “little monarchs” that Hobbes wrote about disappeared into the armed forces of big ones.⁵² Fustel de Coulanges claimed that the conquest of the ancient world by Rome was facilitated by the fact that aristocratic leaders in many city-states thought that submission to Rome would protect them from popular forces at home (1956, 373–74). And the recent conquest of much of Afghanistan by the organization known as the Taliban was made possible in part by the fact that leaders of opposing groups expected to profit from submitting to it (Rashid 2000, 35).⁵³

If giving up their political independence is very costly for the leaders of states, however, then only relatively weak states would have to agree to do it. Thus equality of power among individual states will lead to system stability, but system stability will not require “balancing.”

Waltz made an influential distinction between “internal” and “external” balancing, which seems to imply that they are just two ways of doing the same thing (1979, 118). But according to Waltz, internal balancing consists of “moves to increase economic capability, to increase military strength, to develop clever strategies.” Since the distribution of such capabilities among all states will influence the distribution of any goods to be divided, all states can be expected to be interested in strengthening their own capabilities relative to others’. If no state has a natural advantage over all the others, the result of such competitive efforts might well be that states are and remain relatively evenly matched, just as one team need not dominate the National Football League forever even if there were no rules whose purpose is to avoid the creation of “dynasties.”

If so, and if the leaders of all states also place a high value on remaining independent, then this will lead to agreements among them that pre-

52. See, for example, Henry Kamen’s (2003) recent account of how the Habsburg family used the scanty resources of the Spanish monarchy to organize predators from all over Europe to share in the benefits of a global empire. For an analysis of the organization created by entrepreneurial Spanish monarchs, see Glete 2002, 67–139.

53. Formal models of interstate war usually assume that the size of the object in dispute is fixed and an actor that loses its independence loses everything. But this clearly need not be true. Waltz himself said that “the system won’t work if all states lose interest in preserving themselves. It will, however, continue to work if some states do, while others do not, choose to lose their political identities, say, through amalgamation” (1979, 118).

serve their independence, even though wars may be necessary to reveal the true distribution of power. But this does not imply that states have engaged in “external balancing,” if that consists of joining with weak states against stronger ones or forgoing the opportunity to absorb defeated states when it arises.

It is not surprising to discover that this is Clausewitz’s own explanation of the ability of European states to maintain their independence. The reason “even gifted commanders and monarchs . . . had to be content with moderate success,” he wrote, “lies with the balance of power in Europe.” Political relations among European states, he said,

had become so sensitive a nexus that no cannon could be fired in Europe without every government feeling its interest affected. Hence a new Alexander needed more than his own sharp sword: he required a pen as well. Even so, his conquests rarely amounted to much. (Clausewitz 1976, 590)

And the reason conquests rarely amounted to much was that the military resources available to states were limited and commonly known:

Their means of waging war came to consist of the money in their coffers and of such idle vagabonds as they could lay their hands on either at home or abroad. In consequence the means they had available were fairly well defined, and each could gauge the other side’s potential in terms both of numbers and of time. War was thus deprived of its most dangerous feature—its tendency toward the extreme, and of the whole chain of unknown possibilities which would follow. . . .

The conduct of war thus became a true game, in which the cards were dealt by time and by accident. In its effect it was a somewhat stronger form of diplomacy . . . in which battles and sieges were the principal notes exchanged. Even the most ambitious ruler had no greater aims than to gain a number of advantages that could be exploited at the peace conference. (589–90)

The French Revolution and Napoleon removed the limits on France’s military resources, Clausewitz wrote, and made them more difficult to measure. However, the other states of Europe were able to recover from their surprise before it was too late, and therefore even “the terrible *Bonaparte*” was unsuccessful. Nonetheless, we should note, France did not lose its independence.

Note that there are two elements to Clausewitz’s explanation of interstate conflict in Europe prior to Napoleon: the fact that the military

resources available to states were limited and the fact that they were commonly known, or, as he put it, that “each could gauge the other side’s potential”:

The enemy’s cash resources, his treasury and his credit, were all approximately known; so was the size of his fighting forces. No great expansion was feasible at the outbreak of war. (590)

One possible interpretation of what he wrote is that it was the limits that were important:

Knowing the limits of the enemy’s strength, men knew they were reasonably safe from total ruin; and being aware of their own limitations, they were compelled to restrict their own aims in turn. (590)

One might infer from these comments that states were able to maintain their independence because no one had the ability to threaten it.

However, Clausewitz denied this:

Even a royal commander had to use his army with a minimum of risk. If the army was pulverized, he could not raise another, and behind the army there was nothing. That enjoined the greatest prudence in all operations. Only if a decisive advantage seemed possible could the precious instrument be used, and to bring things to that point was a feat of the highest generalship. (590)

Thus absolute war even in the eighteenth century was risky, and what enabled states to minimize the risk of complete defeat was the fact that their capabilities were commonly known, which enabled them to fight limited rather than absolute wars. Our discussion of the relation between bargaining and war helps explain why this might be true.

By removing many of the limits on the capabilities of eighteenth-century states (which, as Clausewitz said, “in a sense consist only in man’s ignorance of what is possible”), the military revolution that the political revolution in France made possible made it more difficult for states to have a common understanding of what their relative capabilities were. As a result, in the wars with Napoleon,

There seemed no end to the resources mobilized; all limits disappeared in the vigor and enthusiasm shown by governments and their subjects. . . . The sole aim of war was to overthrow the oppo-

nent. Not until he was prostrate was it considered possible to pause and try to reconcile the opposing interests.

“Will this always be the case in the future?” he asked. “From now on will every war in Europe be waged with the full resources of the state, and therefore have to be fought only over major issues that affect the people?” Clausewitz declared himself unable to answer this question (593).

The French Revolution compelled the predatory rulers of Prussia, Austria, and Russia to make war “a concern of the people,” as Clausewitz put it, which is consistent with Kant’s claims about the effect of recurring wars on the constitution of states (592). Kant’s answer to Clausewitz’s question about what the effect of this change would be was optimistic. The nineteenth century seemed to support such optimism, but the twentieth century did not. In deciding whether Kant may be right in the longer run, we must bear in mind the distinction Clausewitz made between the magnitude of states’ capabilities and their ability to estimate them consistently.⁵⁴

Bargaining and the Recurrence of War

Before discussing the effect of the recurrence of war, let us consider what we have learned about how to explain it. Blainey’s explanation for the recurrence of war during the eighteenth century was that those wars were indecisive and indecisive wars tend to produce short periods of peace. The reason the Napoleonic wars led to a lasting peace, he claimed, was that they were decisive (Blainey 1988, 112–13).

Blainey thought that indecisive wars produced short periods of peace because peace required “a clear ladder of international power,” which indecisive wars did not establish (1988, 109). As we saw, his justification for that proposition was that wars end when one state shows that it is clearly stronger than the other. But if eighteenth-century wars were indecisive they must have ended even though no state had done that. Thus eighteenth-century wars are not explained by Blainey’s thesis; they are counterexamples to it. Blainey claimed Clausewitz’s support for his idea, but, as we have just seen, Clausewitz’s explanation of the limited nature of eighteenth-century wars was that statesmen found it *easy* to arrive at a mutual understanding of their military capabilities, and his explanation

54. For a discussion of Clausewitz’s own ideas on the relation between war and the development of the European state and his involvement in the reform movement in Prussia stimulated by the wars with Napoleon, see Paret 1985.

for the “decisive” nature of the Napoleonic wars was that this ability was upset by the French Revolution.⁵⁵

The analysis of the relation between bargaining and war offered previously implies that two conditions must be satisfied for a ruler to try to overturn a prior peace settlement by going to war: (1) there must be some change that leads him to think that his bargaining power has increased by more than another ruler (or rulers) believe, and (2) he must place a higher value on a military contest that might reveal his true military capabilities than on the terms of the existing settlement. What Clausewitz said about eighteenth-century warfare implies that rulers considered it a safe and inexpensive way of revealing small changes in their relative bargaining power and that this was the result of the fact that they had a good common understanding both of their relative military capabilities and their reluctance to take large risks. Thus they all felt free to challenge prior agreements “as soon as a change of circumstances shall have given fresh strength to the claimants,” as Saint-Pierre said.⁵⁶ Since rulers in the eighteenth century were engaged in constant attempts to engineer changes in their circumstances, relatively frequent but limited wars were to be expected.

How things might have been different after Napoleon’s wars is not so clear. But Clausewitz’s discussion of the recurrence of war in the eighteenth century reinforces the importance of our earlier observation that the ability to bargain while fighting, while it may reduce the severity of war, may also increase its frequency.

What Next?

Throughout this discussion I have assumed that land is valued by rulers only for its contribution to their wealth or the wealth of their extended families or their followers. But land can be a source of military capabilities as well, and therefore the redistribution of land might lead to the redistribution of military capabilities and thus a change in the expected value of a military contest. This, of course, was true of Europe in the eighteenth century, and it is what the literature on the balance of power has always assumed.⁵⁷ Thus our analysis of warfare among predatory rulers is seriously incomplete.

55. For an interesting argument that it is the length of wars and not their decisiveness that determines how much information they reveal, see Smith and Stam 2002.

56. See the discussion of Saint-Pierre’s ideas in chapter 2.

57. This can be confirmed by even a cursory reading of any standard work on this subject, for example, Gulick 1955.

Moreover, I noted earlier that in thinking about the relation between bargaining and war one could not, as most of the literature on bargaining does, ignore the question of how agreements are to be enforced. Yet so far I have ignored it. When there is a connection between the object in dispute and the relative bargaining power of the adversaries, this problem is especially complex. I will try to correct both these deficiencies in the next chapter.

Enforcing Agreements

As we have seen, Hobbes's statement that "covenants, without the sword, are but words, and of no strength to secure a man at all," along with the Prisoner's Dilemma game, have helped make plausible the view that the central difference between international politics and domestic politics (or between "anarchy" and "hierarchy," to use Kenneth Waltz's terminology) is that within states contracts are enforceable and among states they are not. But this view rests on two confusions. One is a confusion about the relation between the enforceability of contracts and conflict over their terms, and the other is a confusion about enforcement.

The confusion about the relation between enforceability and conflict is illustrated by the comparison discussed in the previous chapter between wars and strikes. The fact that agreements that end strikes are enforceable does not prevent strikes but rather encourages them, since it increases their value. Similarly, we saw that wars might be more frequent if negotiated settlements are possible than if they are not.¹

Of course, if a contract signed after a strike could never be altered, then no further strikes would be possible. But it would probably be harder to reach agreement on the terms of a permanent contract than one with a limited duration and harder to enforce it were it to be signed. A more promising way to deal with the cost of strikes is to try to reduce the inefficiency of bargaining over the terms of a contract or to agree on an alternative way of resolving conflicts, such as compulsory arbitration.

1. This is a point that is insufficiently appreciated by writers who argue that interstate norms and international law are more efficacious than is commonly appreciated by Realists and infer that they are therefore more likely to prevent interstate wars than Realists believe (see, e.g., Kratochwil 1989). Hans Morgenthau wrote that "during the four hundred years of its existence international law has in most instances been scrupulously observed" (1948, 211). A willingness to fight over the rules that are to govern the international order may be the result of a belief in their efficacy, and the "anarchic" nature of the international order is in part the result of attempts by states to avoid conflicts by limiting the extent to which they will be constrained by that order. (On this point, see the section of Waltz's book *Theory of International Politics* called "The Virtues of Anarchy" [1979, 111–14].) And a rule that says effective control over a well-defined tract of territory is sufficient to elicit recognition as a sovereign state by other rulers greatly increases the value of capturing that territory by forceful means.

Thus a better way of characterizing the difference between domestic and international politics might be that within states there is agreement on procedures for resolving conflicts that are more efficient than the unrestrained bargaining that takes place among them. The important question, then, is how costly interstate bargaining must be and whether its costs could be reduced.

The confusion about enforcement derives from the idea, expressed in Hobbes's famous statement just quoted, that enforcement of agreements (including agreements to abide by procedures for resolving conflicts) always requires confronting violators with superior force. As we saw in chapter 3, such reasoning rests on a confusion between the government's role in enforcing agreements among individuals and the enforcement of the agreements that define the state itself. One such set of agreements defines the organization of the government, another its relation to its subjects, and a third the boundary between its territory and the territory of other states, and they are all subject to renegotiation by the use of force. There is no external enforcer of any of them, and therefore what enforces them all is a comparison of the benefits they provide with expectations about the consequences of trying to renegotiate them. A more economical way of saying the same thing is that all these agreements must be self-enforcing.

In evaluating any such agreement, therefore, whether hypothetical or actual, a party to it must consider not only its provisions but also how long it could be expected to last and what the consequences might be of an attempt to renegotiate it. Even if one did not expect to gain from future renegotiation oneself, someone else's attempt to renegotiate it would at least be costly and could perhaps lead to another agreement with different terms.

Since force can be used to negotiate the terms of all these agreements, the terms of any agreement that is accepted will reflect both how people are organized for the use of force and the distribution of instruments of violence among them. Changes in either of these conditions can therefore provide the opportunity for renegotiation. Such changes can be caused by exogenous factors, which might be expected or unexpected; they can be the result of efforts made by parties to the agreements to change them; and they can even be a consequence of the agreement itself.

Unexpected exogenous changes can lead to unexpected forceful renegotiation, but if they are totally unexpected they can have no effect on whether agreements can be reached or their terms. Changes that are anticipated with some probability, however, whether exogenous or not, make the negotiation of mutually beneficial agreements even more difficult than the analysis of forceful bargaining in the previous chapter implies.

We will see that it is the possibility of changes in the distribution of

bargaining power that makes the security dilemma seem plausible as an explanation of war. However, the problems caused by expectations of such changes are neither a necessary consequence of the absence of government nor always eliminated by the creation of one.²

The Struggle for Power

We saw in the previous chapter that contests in forcible disarmament are contests to determine the relative bargaining power of the antagonists in subsequent forceful bargaining. Complete disarmament (which Clausewitz compared to pinning a wrestler to the mat) not only eliminates the defeated state's ability to threaten the victor's military forces but also enables the victor's military forces to replace the enemy's police and make take-it-or-leave-it demands of individual members of its government or other residents of its territory.

However, destroying the enemy's army may not eliminate entirely his capacity for organized resistance, perhaps through guerrilla warfare, or his ability to bargain collectively over the terms of a peace settlement. Thus the consequences of victory in a contest in disarmament are not uniform, and inconsistent expectations about what they will be are an additional possible impediment to a negotiated settlement.

In the previous chapter I assumed that victory in a contest in disarmament enabled the victor to appropriate the valuable territory of the loser. Using this assumption I showed that Clausewitz was right in suggesting that there will often be divisions of the prize that the antagonists will prefer to fighting over all of it, and any division that is agreed to will be influenced by the distribution of military capabilities between or among them at the time the agreement is made, whether this happens prior to fighting or after fighting has begun.

But contests in disarmament can be preceded by contests in armament. Like contests in disarmament, contests in armament can be inefficient. However, the source of this inefficiency is often misunderstood. Moreover, Kant's ideas suggest that such contests can have beneficial effects that are easy to overlook.

One obvious form such a contest might take is what is commonly called an arms race. Arms races might be inefficient because their result can be that the adversaries' relative military capabilities remain the same but they both spend more on arms. If so, they would both be better off if they could agree to reduce their military forces in a way that did not alter

2. The following discussion owes a great deal to Fearon 1994 and 1995b. For a general theoretical analysis of the effect of expected future changes in bargaining power on the ability of adversaries to avoid conflict, see Powell 2004.

their relative military capabilities. However, if one complied with such an agreement and the other did not, then the one that did not would gain a military advantage. As we saw in chapter 1, this might appear to imply that arms races exemplify the inefficiency explained by the Prisoner's Dilemma game and therefore that the explanation for them is the absence of any means of enforcing an arms limitation agreement.

But this inference is wrong, because it overlooks the possibility that one state's maintenance of an arms limitation agreement can be made conditional on the other's behavior, in which case violating the agreement would be self-defeating since it would lead to the rearmament of the other side. That is why arms limitation agreements are accompanied by arrangements for detecting cheating. Of course, the costs associated with such arrangements, and whatever residual arms capability is necessary to provide a base for rearmament if necessary, are themselves inefficient, but the magnitude of the inefficiency is far less than the inefficiency that an unrestricted arms race would entail.

However, such an agreement is possible only if states can agree on some level of armaments that they want to maintain. But if any distribution of military capabilities determines the terms of an agreement that might be accepted as an alternative to a contest in forcible disarmament, then it will not be possible to agree on some distribution of military capabilities without agreement on the terms of a political settlement that both would prefer to war—in the situation analyzed in the previous chapter, that would be an agreement on the distribution of territory. Thus just as states may fight if they have inconsistent expectations about the terms of an agreement that would be accepted after fighting, so may they engage in an arms race if they have inconsistent expectations about the terms of an agreement that would be accepted after a competition in armament.

But in that case an arms race resembles a limited war: it is a contest conducted prior to a contest in forcible disarmament, whose outcome will determine expectations about that contest. And like a limited war, an arms race may reveal enough information about the states' relative bargaining power that a more costly and dangerous conflict can be avoided—recall Clausewitz's observation that it was lack of information about the level of armaments that Napoleon could mobilize that accounted for the length and the cost of the wars that were fought in the aftermath of the French Revolution. Like any bargaining, therefore, such an arms race would be inefficient. But like limited war, its inefficiency would be far less than the inefficiency associated with a contest in forcible disarmament.

Moreover, just as states may engage in arms races, so may they also engage in competition in creating the resources from which armaments are derived. But these are bureaucracies, modern economies, tax systems, and the ability to persuade the bulk of one's male population to fight. As we

saw in chapter 2, it was this competition that drove much of the process of state building in Europe.³ Rousseau was of two minds about whether on balance the effects of this competition were good or bad. Kant thought that in the long run they would be good, but during Hitler's and Stalin's time this might have seemed naive. The end of the cold war and the liberalization of China made many people take Kant's ideas more seriously.

During the winter, the future looks bleak, but during the spring optimism revives. To know what to expect, we must know what explains the cycle of the seasons. Unfortunately, we do not know what to expect from the struggle for power. But its explanation is neither that men are endowed with "a perpetual and restless desire of power after power, that ceaseth only in death," nor that "covenants, without the sword, are but words, and of no strength to secure a man at all." Three other explanations are possible: (1) the participants all remain too optimistic about how they will fare if the contest is continued to agree on the terms of an agreement that might end it; (2) it is not possible to coordinate the actions of everyone whose cooperation would be necessary to end it; and (3) an agreement to end it would not be self-enforcing.

Whatever the explanation, this competition is one important source of the recurring changes in bargaining power that make more difficult the construction of other stable, self-enforcing agreements. One of its possible consequences is the creation of incentives for states to attack their adversaries before they are attacked by them.

Incentives to Attack First

One of the recurring controversies about arms races is whether they make war more likely. I have just pointed out one way an arms race could make war less likely, though in any given case this effect might be concealed by the fact that it ended in war: it may have failed to reveal enough information about the adversaries' relative bargaining power to make an agreement possible without fighting. However, an arms race could also make war more likely, if it leads to a military advantage from attacking before one's adversary does.

In any arms race one side or the other may achieve a temporary advantage. But the one that does may not want to attack immediately to capitalize on it, since it may be optimistic that the arms race will eventually reveal enough information about the two sides' military capabilities and interests to lead to a negotiated settlement on favorable terms. If it believes that its adversary has a long-term advantage in the struggle for power,

3. For a recent survey, see Glete 2002.

however, it may have an incentive to attack before its own temporary advantage disappears. And whatever the distribution of military capabilities, it is possible that there is an advantage to being the one that attacks first. The first situation may lead to a preventive war, and the second to a preemptive one. I will discuss preemptive wars first. Throughout the discussion I will continue to assume that the competitors are predatory rulers competing over the distribution of valuable territory.

Preemptive Wars

The possibility of a preemptive war became a preoccupation during the cold war, when many people worried that an incentive to attack first could lead to a nuclear war even though neither side wanted to fight one. Many discussions of this possibility are based on the assumption that a state can only choose between attacking another state and not attacking. In those circumstances, a state that chooses not to attack exposes itself to the possibility of an attack by the other state. If it believes the other is about to attack, its choice is then not between war and the status quo but between two different wars: a war in which the other state attacks first and a war that begins with its own attack. Thus if a state mistakenly believes that a nuclear war is inevitable but prefers one in which it attacks first, it appears that two states could fight a nuclear war even though neither in fact preferred such a war to the status quo. Such a war might be called “unwanted” or even “inadvertent,” though it would have been intentional.⁴

But this reasoning ignores the possibility of a negotiated settlement as an alternative to war. If a negotiated settlement is possible, the effect of incentives to attack first becomes more complicated.

In discussing the balance of power in the previous chapter, I pointed out that it can be thought of in two ways—as the distribution of military capabilities or potential or as the probability with which each side in a military contest could be expected to disarm the other—and I assumed that only the former was relevant to estimating the latter. But if there is an advantage to being the first to attack, then the probability of victory depends not just on the distribution of military capabilities but also on who attacks first.

4. One must be careful with words like *accidental* or *inadvertent* when applied to wars. Weapons can be fired or detonated accidentally (with very costly consequences if they are nuclear weapons), or airplanes can accidentally fly over another state's airspace, but it is not obvious how a military contest could occur by chance. However, an accident could cause a war that might be called inadvertent or unintended if it caused a state with an incentive to attack first to believe that its enemy was about to attack. For a recent discussion and citations to the literature, see Powell 2003.

We saw in the previous chapter that any state dissatisfied with the existing distribution of territory would have to attack the other state if it wanted to compel a redistribution of it, and therefore the probabilities of success that determine the range of possible settlements would have to take that into account. If there is an advantage to attacking first, then each side's probability of success would have to be its probability of disarming the other, conditional on its being the first to attack. In that case the probabilities in figure 6 need not add up to one, and so the effect is similar to the effect of inconsistent expectations discussed in the previous chapter. The result could be to narrow the range of possible agreements (if there is an advantage to attacking) or to expand it (if there is a disadvantage to being the attacker). But even if there is an advantage to attacking, this need not eliminate the range of possible agreements entirely. And if it does not, then the advantage to attacking first will already have been taken into account in determining the existing distribution of territory.

Thus if an advantage to attacking first is to lead to an attempt to change the status quo, the advantage must have increased. But a state confident of having the advantage of attacking first might still prefer a compromise settlement to fighting. And therefore if the effect of attacking first is common knowledge, and a state chooses to attack rather than to demand a concession, this cannot be just because there is an advantage to attacking first but must be because there is an advantage to a *surprise* attack. In that case it could not demand a concession without revealing its intention to attack, and the other state could not commit itself not to take advantage of that information. And therefore the state optimistic about capturing the advantage of attacking first might attack and then demand a concession, after it had secured its advantage.⁵

If it did, however, the resulting war would be no more unwanted or inadvertent than any other war, since the attacking state would expect to gain a bargaining advantage by attacking. It would, however, have been inefficient if the victim had been prepared to offer a concession that the attacker preferred to fighting.

Moreover, if the only information that is required to reach an agreement is information about which side owns the advantage of attacking first, then it should be possible to reach an agreement soon after fighting starts, and therefore wars fought solely because of first strike advantages should be short. Indeed, this is one possible explanation of some of the "limited aims" wars that Clausewitz wrote about, since the advantage of attacking first might consist of the ability to capture a piece of lightly defended territory before the enemy is able to respond. Ownership of this

5. An incentive to attack by surprise is therefore similar to incentives to conceal other components of a state's military capabilities, in order to prevent the enemy from taking counter-measures against them.

territory could then be confirmed in a peace agreement, or it could be traded for some other gain instead.⁶

The preemptive war scenario assumes not just that a state sees an advantage to attacking without warning but also that its potential victim mistakenly anticipates an attack and decides to try to attack before the first state is able to. But if the incentive to attack without warning is the result of being able to catch the victim unawares before he is able to mount a proper defense, the potential victim might not want to preempt an attack but prefer instead to defend himself against it, thereby nullifying the advantage of a surprise attack. For preemption to be considered, the optimal military response to an expected surprise attack must instead be a surprise attack of one's own, which is not implied merely by the existence of an advantage from attacking first.

Moreover, this scenario overlooks the possibility that the second state might offer a concession instead. A state that had the advantage of a surprise attack would have to give it up if it demanded a concession from the victim. But a state that expected to be the victim of such an attack might prefer to concede the advantage of attacking first to its adversary, since if it attacked instead it could not be certain of forestalling the enemy's attack, and even if it did it would still have to face the cost and risk of fighting. If so, then the result might be an unnecessary concession, but not a war that neither wanted to fight.

Note that this is especially likely to be true if the war to be fought were a nuclear war, and therefore it is actually easier to construct a scenario leading to an inadvertent conventional war than a nuclear one.⁷

Note also that, if neither state is optimistic about the outcome of a competition for the advantage of a surprise attack, it may be possible to agree on measures that would reduce or eliminate it.

Preventive Wars

A preemptive war would be the result of a state's attempt to prevent an adversary from acquiring the bargaining advantage that a surprise attack would give it. A preventive war would be the result of a state's attempt to

6. World War I is sometimes explained as the result of Germany's incentive to attack first to avoid fighting Russia and France simultaneously. But such an explanation fails to explain why the war did not end as soon as it was clear what advantage Germany had gained by attacking first. For a discussion of why World War I lasted as long as it did, see Goemans 2000. Of course, the fact that the war did not last long would be small consolation if it were a nuclear war.

7. In the case of nuclear war, one must also explain why there would be an advantage to attacking first and show how a state could come to be confident that another was about to attack if it were common knowledge that the only reason either would ever attack was to preempt an expected attack by the other side.

prevent an adversary from acquiring the advantage that an expected future increase in its military capabilities would give it.⁸

Like preemptive wars, preventive wars are often discussed as though states faced a simple choice between attacking now and waiting to fight a less desirable war later. If so, then the difference between them is that in a preemptive war the disadvantage to waiting is that the enemy gets to attack first and its attack is imminent, while in a preventive war the disadvantage to waiting is that the enemy will grow stronger with time but the attack will come later. Once again, however, the problem is more complicated if one considers the possibility of a negotiated settlement as an alternative to fighting.

We saw that, if a negotiated settlement is possible, the alternative to a preemptive attack is a preemptive concession. The alternative to a preventive war, however, is acceptance of the possibility of having to make a concession in the future. But a concession will be required only if (1) the adversary's military capabilities increase as expected, (2) it is not possible to compensate for that increase by actions to increase one's own capabilities, and (3) the adversary's preferences will enable it to translate its new relative military capabilities into sufficient bargaining power to compel a concession.⁹ These conditions might imply that the cost of the future concession should be discounted heavily.

If the only barrier to agreement without fighting is knowledge of which side will have the advantage of attacking first, a preemptive war, we saw, might be expected to be short. The aim of a preventive war, however, would not be to compel an immediate agreement that reflected the current distribution of military capabilities but to forestall a future agreement when the distribution of military capabilities would be less favorable. But that might imply that the initiator of a preventive war, if the war went well, could not cash in his success by accepting an early negotiated settlement but would have to proceed until he had weakened his adversary to the point that he no longer feared its future military capabilities—though he might quickly accept a negotiated settlement if he soon became pessimistic about how the war would end. A successful preventive war might therefore be expected to be far longer and more costly than a preemptive one.¹⁰

Thus the possibility of a preemptive war entails a choice between a

8. Note that the expected future increase in the enemy's military capabilities might consist of a future advantage from attacking first, a fact that makes it even easier to confuse the two problems. For an influential early discussion of preventive war, see Levy 1987. A more recent discussion is in Fearon 1995b. Preventive war is the main theme of Copeland 2000.

9. Recall the discussion in the previous chapter of the determinants of bargaining power when the alternative to agreement is a contest in disarmament.

10. This would not be true if the attacker only wanted to destroy a part of the enemy's military capabilities, for example, a nascent nuclear weapons capability.

war that could be expected to be short (though one in which one might not be certain of having the advantage of attacking first) and an immediate concession, while a preventive war entails a choice between a longer and more costly war and the uncertain prospect of a possible future concession. The explanation of both is the inability of states to commit themselves not to exploit a future change in their relative military capabilities, and the effect of both is to reduce the range of agreements that are feasible as alternatives to war, but not necessarily to eliminate it. It seems plausible that in a world of conventional weapons, at any rate, fluctuations in relative military capabilities will lead more often to contemplation of preventive than preemptive wars; but the conditions that must be satisfied for a preventive war to be chosen imply that they will often be rejected.¹¹

The Security Dilemma Reconsidered

While any war is inefficient if the negotiated settlement that ends it would have been preferred by both combatants to fighting, many people came to believe that Herz's security dilemma implied that wars could occur even though both sides actually preferred the prewar status quo. We saw that that is not true. However, it can be true of both preemptive and preventive wars (though it need not be): the state whose attack is preempted might actually not have intended to attack, and the state whose increase in power is prevented might never have challenged the status quo.¹²

The absence of government alone does not imply that either preemptive or preventive wars will occur, so they cannot be explained solely by "anarchy." But might governments nonetheless prevent them, so that the absence of government is at least a necessary part of the explanation for their occurrence?

The belief that governments can reliably prevent such conflicts is another example of the confusion between the role of governments in regulating conflicts among the people they govern and the role of governments as parties to conflicts with the people they govern. The organizational advantage that governments have over potential domestic opponents or that political leaders have over potential dissidents within the government implies that leaders of opposition groups will often have an incentive to launch a coup d'état or a rebellion before the government

11. Bismarck famously said that "preventive war is like suicide from fear of death" (quoted in Levy 1987, 103). Of course, people do commit suicide from fear of death.

12. See the discussion of the security dilemma in chapter 1. Note that for both states to prefer the status quo to war, it is only necessary that, given the existing distribution of military capabilities, both would prefer the existing distribution of territory to the expected outcome of trying to change it. Thus even predatory rulers competing for valuable territory might both be satisfied with the status quo.

is able to act against them or acquires the power to do so. And for that reason, many acts of political repression are designed to preempt or prevent dissent rather than to respond to it. Incentives to attack first can add to the inefficiency of forceful bargaining wherever it occurs, and it occurs not only among governments but also within them and between governments and the people they govern.¹³

Since the end of the cold war, some scholars have used the security dilemma to explain the civil wars in the Balkans and the attempted genocide in Rwanda. In doing so there has been a tendency to equate the security dilemma with an incentive to wage a preventive war. And since everyone assumes that the security dilemma is caused by anarchy, there is also a tendency to assume that if the cause of civil war is a security dilemma it must be because the government collapsed and plunged everyone into anarchy.¹⁴ But the security dilemma does not imply the occurrence of preventive wars, and governments can lead to preventive civil wars that would not have occurred had the antagonists been separate states with separate military forces—hence the possibility of resolving domestic conflicts by partition.

Offense and Defense Reconsidered

To defend something is to “make or keep [it] safe from danger, attack, or harm” (*American Heritage Dictionary*). Thus the aim of the defense is to maintain the status quo, and the aim of the offense is to change it. In the context of contests in disarmament, if force is not used the status quo will be unchanged, so the offense must be the initiator of the use of force. And if the use of force is to have any effect, the initiator must engage the military forces of the enemy, wherever they are. And therefore, given an equal distribution of military capabilities, to say that the offense has an advantage could mean either that the initiator of the contest is more likely to win a contest in disarmament than the side that awaits an attack or that the side that fights on its home ground is more likely to lose.¹⁵

The fact that incentives to attack first can make a connection between the security dilemma and war seem plausible is perhaps one reason why many people have found persuasive Robert Jervis’s (1978) claim that the severity of the security dilemma depends on whether the offense or the defense has an advantage in military contests, since saying that the offense has an advantage over the defense clearly implies that, other things being

13. See Fearon’s (1994) discussion of ethnic war in the Balkans and Weingast’s (1998) discussion of the U.S. Civil War. See also the discussion in Lake 2003.

14. See the essays collected in Walter and Snyder 1999, especially Snyder and Jervis 1999.

15. See Clausewitz’s discussion of the distinction between offense and defense in Clausewitz 1976, 357–59.

equal, the attacker has an advantage in contests in disarmament.¹⁶ However, we saw that, while the existence of an advantage to being the attacker may reduce the range of feasible agreements, it need not eliminate it. And if it did eliminate it, then *it could not be true that both states preferred the status quo to war*. Moreover, if there is a range of agreements that both states prefer to war, then the fact that the attacker has an advantage will lead to war only if (1) there is also an advantage to attacking without warning and (2) the defender is successfully surprised or the defender's optimal response to an expected attack is a surprise attack of its own.

The connection between the offense-defense balance and preventive war is less direct. However, if the offense has an advantage, then a state would find it more difficult to counter an expected future increase in its adversary's military capabilities than it otherwise would be and would be more optimistic about disarming its adversary if it attacked before the increase occurred. And therefore, while it is not necessary that attackers have an advantage for preventive wars to occur, they might be more likely if that were true.

But even if offensive advantages make preemptive and preventive wars more likely than they otherwise would be, it does not follow that they make wars more likely. We saw in the previous chapter that wars might be frequent even if the probability of success in a contest in disarmament was completely unaffected by whether a state attacked first or not, since it is not necessary to disarm one's adversary in order to use military force to change the status quo—force can be used to extract a concession instead. And, as we learned from Clausewitz, the less likely it is that a state will be disarmed in a war, the more attractive war becomes as a form of coercive diplomacy. In those circumstances what determines the frequency of war is not the offense-defense balance but the frequency of changes in relative military capabilities great enough to support an attempt to renegotiate the agreement that ended the previous war.¹⁷

Bargaining over the Distribution of Power

So far I have assumed that what is at issue is the distribution of valuable territory among predatory rulers and argued that the distribution of mili-

16. See the discussion in chapter 1. It is also easy to be misled and think that if the offense has an advantage over the defense then the offense is likely to be successful, which need not be true.

17. Note that Clausewitz claimed that "defense is the stronger form of waging war" but did not think that was inconsistent with the frequency of wars in the eighteenth century (1976, 359).

tary capabilities, by influencing the distribution of bargaining power, determines the distribution of valuable territory.

But this overlooks one of the most important facts about interstate conflicts: the distribution of territory can affect the distribution of military capabilities. Territory can affect military capabilities in two ways: its location or topography can directly affect the ability of states to deploy military forces against each other, and population and economic resources located within it can be converted into military capabilities. Thus an agreement about the distribution of territory between states might influence the subsequent distribution of military capabilities between them and therefore change the relative bargaining power on the basis of which the agreement was reached. A concession of territory by state A to state B, therefore, could enable state B to demand a further concession at a later point, and so on, until state A had ceased to exist. Since no state could commit itself not to take advantage of such an increase in its future bargaining power, it would appear that no agreement could be self-enforcing and therefore none was possible.

It would be wrong to leap to the conclusion that every possible agreement that states might accept instead of fighting has this property. Nonetheless, many obviously do, and it clearly implies a major constraint on their ability to reach agreements without fighting. It is what makes describing a struggle for territory as just a struggle for power seem plausible, and, as we will see, it helps explain much that has been written about the balance of power. In spite of that fact, it has not received much analytical attention.¹⁸ Let us see if we can figure out what its implications might be.

One obvious implication is that territory with this property becomes even more valuable, since it not only has value in itself but also makes the forceful acquisition of additional valuable territory more likely. And territory that has no economic value might become valuable because of its strategic significance.

The problem, however, is not that territory might be valuable enough to be worth fighting for, since we have already assumed that to be true. Even if territory were just the same thing as military power, for example, then, in figure 6, p would always equal q and only rulers whose risk acceptance was great enough to outweigh the expected costs of an absolute war would ever want to fight for more. And if rulers are risk averse, there

18. One of the few discussions is in Fearon 1995c. For a contemporary example of the problem, think of the claim made by some Israelis that trading land for peace with the Palestinians would be self-defeating, since giving the Palestinians land would enable them to demand more.

might be a wide band of territorial distributions around q that would be preferred by both sides to a military contest for all of it.

The problem is rather that an agreement that changed the relation between the distribution of territory and the distribution of power would not be self-enforcing. Since, as we saw, the opportunity to coerce a territorial concession by fighting makes war more likely than it otherwise would be, the fact that states could not be expected to make a concession to avoid an absolute war might actually make war less frequent, not more. But this increased stability in the distribution of territory would have a price: if the disparity between the distribution of territory and the distribution of power were great enough that one state or another would prefer a contest for all of it to the existing distribution, the contest could not be prevented by a negotiated settlement.

However, such an agreement would not even be considered unless there were an initial disparity between the distribution of power and the distribution of territory, which would imply that the two cannot simply be equated with each other. One reason such a disparity might exist is that states differ in their ability or willingness to convert the economic resources available in the territory they control into military capabilities.¹⁹ And since territory that a state concedes to an adversary as an alternative to war might not immediately increase its adversary's military capabilities, the state making the concession might hope to compensate for it by increasing the resources it mobilized from its remaining territory in the meantime. If so, then making the concession might be preferable to fighting.²⁰

This is what came to be known as "appeasement" in the period preceding World War II. One of the main arguments against appeasement is that concessions strengthen one's adversary, enabling him to demand more later. But a possible rebuttal is that appeasement buys time that can be used to increase one's own military capabilities.²¹

Moreover, even if appeasement is rejected and war occurs, the result-

19. Another obvious possible reason for such a disparity is that not all territory has equal military value.

20. James Fearon (1995c) has shown that if states discount the future then there are conditions under which states would prefer to make such a concession even if they did not expect to be able to compensate for it by increasing their own military forces. Note also that since preventive war is the result of an expected future increase in a state's military capabilities, the connection between territory and power could make it possible for a state faced with the prospect of a preventive war to avoid it by conceding in the present some of the resources from which its military capabilities are derived.

21. Thus appeasement of Hitler by Britain prior to World War II has been defended as allowing Britain time to rearm. Another, independent, argument against appeasement is that it may cause the adversary to make incorrect inferences about one's preferences, leading him to expect more concessions later even if the distribution of military capabilities is unchanged.

ing war need not be Clausewitz's absolute war. If territory changed hands as the war progressed or if the defender's success led the attacker to be less optimistic about the ultimate outcome, then at some point the distribution of territory might fall once again within the zone of agreement in figure 6, and the two sides could then agree to accept the status quo rather than continue fighting. This, then, is another possible explanation for the "limited aims wars" that Clausewitz wrote about.

But if the alternative to appeasement is not absolute war but limited war, then a state contemplating appeasement confronts a choice between the loss of territory for certain and a limited contest whose outcome might be that it retains the territory it would have conceded. This clearly makes the rejection of appeasement more attractive, and therefore, once again, reducing the expected severity of war makes its occurrence more likely.

Balance of Power Theory Reconsidered

We saw in the last chapter that the relation between absolute war and negotiated settlements among three states is qualitatively similar to the relation when there are just two: given rulers' preferences between divisions of disputed territory and a costly contest for all of it, the distribution of military capabilities among all three states will determine what division they will all accept (if any) as an alternative to fighting. As we just saw, a connection between the distribution of territory and the distribution of power between two states makes peaceful agreement on a division of the territory more difficult. But its effects are somewhat different when there are more than two states.

If negotiated settlements are not possible, then two states fighting a third must anticipate a subsequent war between them if they succeed in disarming their enemy. If negotiated settlements are possible, however, then two allies negotiating with a third will anticipate a negotiated settlement with each other if they disarm their enemy, which is a more attractive prospect. An agreement leaving them each with half the territory, for example, would be more attractive than the prospect of a subsequent military contest that gave each a 50 percent chance of winning all of it. Thus an absolute war in which two states allied together to disarm the third and then divided its territory between them would be more attractive than an absolute war in which the victorious allies had to fight each other after defeating their enemy.

As I noted in the previous chapter, this is the scenario that is assumed implicitly or explicitly by most writers on the balance of power. It poses the problem of the stability of state systems in its starkest form, since if no state is more powerful than all the others combined, there will always be some coalition of states that is more powerful than an individual state, and

if they could always agree on the division of their victims' territory after defeating it, then it is not clear how any interstate system could be stable.²²

As we saw, one possible answer is that a negotiated settlement among all three states is possible, which might lead to a territorial concession by the third state but not to the elimination of any of them. But a connection between the distribution of territory and the distribution of power would make such negotiated settlements problematic, because they might not be self-enforcing.

However, while this may inhibit agreements among all three states, it need not inhibit an agreement between two victorious allies about how to dispose of the territory of the third state if they disarm it. This is because the two allies would not be redistributing territory they already possessed (which would lead to a change in their relative power) but redistributing territory that belonged to the third state (whose power would be irrelevant if it is eliminated). Thus they could divide the territory of the defeated state in such a way that the prewar distribution of power between them was not altered.

This possibility is illustrated in the following description of predation in Renaissance Italy by Garrett Mattingly:

Historians have been able to discover one general principle in sixteenth-century diplomacy related to the idea of national interest, the principle of the balance of power. There are, indeed, episodes in the period 1494 to 1559 when it looks as if that principle was really being applied, especially when it was a question of the combination of two or more strong states against a weak one. Here the principle requires such a partition of the victim's territories as not to change decisively the strength of any victor in relation to his partners. . . . But since it really means little more than that the biggest dog gets the meatiest bone, and others help themselves in the order of size, it is hard to be sure that the sixteenth century appreciated the full beauty of a balanced system. (1964, 140–41)

The eighteenth-century principle of "reciprocal compensation" can be explained in the same way. According to Gulick, this principle required that "aggrandizement by one power entitled other powers to an equal compensation or, negatively, that the relinquishing of a claim by one power must be followed by a comparable abandonment of a claim by another" (1955, 70–71).²³

But, as the quotation from Mattingly illustrates, it is easier to satisfy

22. This is why William Riker claimed that writers on the balance of power were wrong and state systems were inherently unstable (1962, 160–187).

23. See also Schroeder 1994, 6–7.

this principle in negotiations among victorious allies than in negotiations between them and their victim. This, then, is a possible justification for the common assumption that three states could not reach a negotiated settlement as an alternative to war but two states could do so after defeating the third.²⁴ And it seems to imply that our discussion of the relation between bargaining and war among three states exaggerates the ability of states to maintain their independence. Both temporary appeasement of the aggressors by their victim and limited wars between them would still be possible, but if appeasement is followed by more appeasement then it cannot protect states' independence, and if the gains from limited wars are cumulative then states may eventually be eliminated. Perhaps we should take another look at why writers on the balance of power thought this would not happen.

As we saw, the idea that weak states could band together to defend themselves from stronger ones is not a good reason, since they will not necessarily do it, and even if they did, wars might nonetheless lead to the elimination of states.²⁵ But some writers on the balance of power claim not only that weaker states join together to "balance" stronger ones but also that they design peace settlements to restore a "balance" if they are victorious. Gulick's well-known book on the balance of power, for example, is not just about the formation of a coalition to counter Napoleon's France but also about the attempt to craft a peace settlement that would restore a balance of power after France was defeated, and Gulick claimed that the "necessity of preserving the components of the system may be taken as a corollary of the balance of power" (1955, 73).

This would imply that the independence of states is protected not by balancing but by the unwillingness of states to eliminate other states even if they are able to do so. But why would states forgo the opportunity to exploit a military victory to the fullest?

It might appear that they would behave in this way only if their sole interest were in protecting the territory they already controlled.²⁶ However, in their survey of interstate conflict during the eighteenth century, McKay and Scott say that

Rulers and statesmen strove ceaselessly to increase the power, and therefore the wealth, of their state. State power was everywhere

24. For an alternative interpretation, see the discussion in Powell 1999, 160. Powell's explanation of this behavior is that it just reflects the relative bargaining power of the victorious allies. And it is true, of course, that an unconstrained division of the defeated state's territory between the victorious allies would reflect their relative bargaining power and therefore their relative military capabilities. But this just provides further reason for thinking that the victorious allies would find it easy to agree on a self-enforcing division of the victim's territory.

25. For a recent extended analysis, see Powell 1999, chap. 5.

26. This is what Schweller claims (1994, 1998).

measured in terms of territorial extent and population, which in turn determined revenue and the size of the army. . . . Additional territory was everywhere the aim of policy. (1983, 211)

If a lack of interest in territorial expansion is required for states to be reluctant to deprive other states of their independence, then it is hard to see why they would have been reluctant to do it in the eighteenth century.

A possible answer is that an agreement between two states to divide the territory of a third state between them is not as self-enforcing as I have made it out to be. The distribution of military capabilities, as we have seen, might be a function of the distribution of territory, but it cannot be equated with the distribution of territory, and therefore the relative power of two victorious states might change even if it is not changed by the distribution of the third state's territory between them. And if it does, then the postwar territorial distribution may no longer be stable.

Moreover, redistributing the third state's territory will make both of the victorious allies more optimistic about capturing all the territory than they were before the war. For example, if they were all initially equal, the probability with which each could expect to capture all the territory might be one-third. If two then join together to eliminate the third and divide its territory in such a way as to leave them still equal, then the probability with which each might expect to capture all the territory will have increased to one-half. Any further increase in the capabilities of one of them will leave it far more optimistic about complete victory than it was initially.

Thus when there are more than two states, one state cannot expand without giving some other state the opportunity to become more powerful. As we saw in chapter 1, this was the basis for Kenneth Waltz's claim that even expansionist states could not simply try to maximize their power. He claimed that this implied that states would form balancing coalitions, but we saw that this was not necessarily true. However, if agreements between victorious allies are not self-enforcing, it may imply that states would refrain from depriving other states of their independence.²⁷

Consider the problem confronted by two equally matched states that have just succeeded in disarming the third. As already noted, if they divide the territory of the defeated state in a way that reflects the current distribution of power between them, then this division may no longer be an equilibrium if a change in their ability to mobilize military capabilities

27. This is the focus of Wagner 1986 and Niou and Ordeshook 1990. However, these works assume that the outcomes of wars are predictable, which makes it hard to explain why they would have to be fought.

leads to a change in that distribution, and in an environment in which such changes occur frequently such an agreement may therefore not last long.

Suppose, however, that they do not redistribute all of the third state's territory but leave it with a reduced amount. If all three states agree that no state should be allowed to become more powerful than the two victors have become, then no state will accept an agreement giving either more territory than it has, and both may prefer to accept the existing distribution to the expected value of trying to defeat the other two.

Moreover, if subsequently one were expected to become stronger, then the other two would have an opportunity to wage a preventive war against it. McKay and Scott wrote of the balance of power:

In practical terms the balance of power meant simply that no one state, or alignment, should become too powerful; and that if it did, the other European states would join together to reduce its power. (1983, 211–12)

It makes little sense to think that two states would attack a third because the third state was too powerful, as this passage suggests, but two states might attack a third in order to prevent the third from becoming too powerful if it were expected to do so. If they did and were successful then they could divide the territory of the formerly powerful state between them in the same way, maintaining the independence of the victim as a way of securing their own possessions against an uncertain future.

Thus there appears to be a close connection between balance of power thinking and the incentive to wage preventive war: preventive wars are fought to protect states from expected future increases in the military capabilities of an antagonist, and maintenance of a balance of power, interpreted in this way, is designed to preserve the ability of states to cooperate in the waging of preventive wars. Immanuel Kant wrote, for example, that “an alarming increase of power” in another state “which has acquired new territories”

is an injury to the less powerful state by the mere fact that the other state, even without offering active offence, is *more powerful*; and any attack upon it is legitimate in the state of nature. On this is based the right to maintain a balance of power among all states which have active contact with one another. (1797, 167; emphasis in original)

An agreement between two successful aggressors that leaves some of the military capabilities of their victim intact is therefore like depositing some

of the military capabilities at their disposal in an escrow account that can be used against either of them, protecting both against the possibility that their agreement might be overturned by a subsequent change in their relative military capabilities.²⁸

But we have already observed that, in a world of three evenly matched states, the expectation that two states, if victorious, could reach an agreement about how to divide up the territory of the third would increase the expected value to them of an attempt to disarm it. And therefore, if “balancing” provides a means of enforcing an agreement that would otherwise be unenforceable, it would not be inconsistent with what Waltz called “bandwagoning” (i.e., states ganging up on other states to deprive them of their territory) but would actually make it possible.²⁹

The extent to which the expectation of balancing could be used to support the forcible redistribution of territory depends on how powerful each state is willing to allow other states to become, since obviously the more territory that must be left with the third state, the less attractive is a military contest to capture the rest. If balancing is the result of a willingness to sacrifice some territory in order to make possession of the remainder more secure, it is like buying insurance, and the amount of insurance rulers will choose to buy will depend on both their attitudes toward risk and the amount of it they believe they are exposed to. But this implies both that no definite answer can be given to the question of how powerful other states should be allowed to become and that the leaders of different states may give different answers to it.

If, then, in response to the description of balance of power policies by McKay and Scott just quoted, we ask how powerful is “too powerful,” the answer would have to be that there is no general answer to this question. Moreover, there is no guarantee that all states will buy enough insurance to protect them from the possibility that a single state will become powerful enough to coerce the others into relinquishing their independence.

Recall that McKay and Scott claimed that states were concerned not just that a single state might become “too powerful” but that an “alignment” might as well. We have seen that in a three-state world an unwillingness to allow another state to become as powerful as the other states combined can be explained as a way of controlling the risks that states are exposed to when the distribution of military capabilities is believed to be unstable. But this implies that when there are more than three states they all must also be concerned about the power of potential two-state coalitions.

28. Compare this to Avner Greif's (1998) argument that warring clans in late medieval Genoa solved the problem of enforcing cooperation between them by creating a third party that would provide a balance of power.

29. For a discussion of balance of power thinking in the eighteenth century that emphasizes its close connection with predatory behavior, see Schroeder 1994.

tions, since a two-state coalition could secure its winnings by eliminating all but one of the other states, which would be expected to “balance” between the two winners. And therefore all states must worry that two states, and not just one, might become optimistic enough about their ability to disarm all the others to try it. Thus the interpretation of balance of power thinking offered here can easily be extended to coalitions in a world of more than three states.

When violent conflicts are about the distribution of territory, then, and the distribution of territory affects the distribution of military capabilities, there appears to be a close connection between the sort of behavior emphasized in the literature on the balance of power and the fact that agreements among states must be self-enforcing. However, these conditions do not always hold. Moreover, “moderation” in the behavior of states (as Gulick called it) could also just be a consequence of the ability of states to reach negotiated settlements when they are all relatively evenly matched. And therefore it is not clear how important this factor is in explaining the historical behavior of states, and it is possible that most of the time international politics is best understood as a complex multiactor bargaining process.³⁰

Extended Deterrence and the Balance of Power

The balance of power reasoning just analyzed also casts some light on what is commonly called the problem of “extended deterrence” and helps explain why it received so much attention during the cold war.

The problem of extended deterrence is the problem of how to make credible one state’s commitment to defend another, which was the main preoccupation of U.S. foreign policy during the cold war.³¹ Thus there are at least three states involved: a defender, a potential aggressor, and a third state that the defender is committed to defending. It is commonly assumed that the defender’s goal is to deter an attack on the client state by the

30. Consider, for example, commentaries that portray recent efforts by Russia and China to thwart the “hegemony” of the United States by acting in concert as examples of modern-day balancing. Such actions are more likely designed to influence the terms of agreements that will be reached by all three states than to reflect the role either Russia or China would play in an all-out war with the United States in the future.

31. During the cold war, the problem of extended deterrence led to the problem of how to make threats to use nuclear weapons in defense of client states credible, but the former problem would have existed even without the latter. Indeed, initially nuclear weapons were seen as a way of solving the problem of extended deterrence, since they made it possible for the United States to devastate the USSR at little cost to itself. This changed when the USSR acquired missiles that could transport nuclear warheads to the United States. For a discussion of extended deterrence with citations to the literature, see Huth 1988. For a discussion of the debates about how to make extended nuclear deterrence credible, see Daalder 1991.

potential aggressor by threatening war against it if it attacks and that avoiding war by making a concession would be unacceptable. What is in doubt is the willingness of the defender to respond in this way, and the problem of extended deterrence is how defenders can credibly reveal their willingness to do it when they have an incentive to bluff. Left implicit in debates about this problem is the assumption that the reason for the defender's interest in preventing an attack is the fact that both a successful attack and any possible concession would increase the power of the aggressor and that both defenders who were bluffing and defenders who were not would have an interest in persuading the potential aggressor that an attack would be unacceptable for this reason.

The analysis offered here explains why these assumptions might be true, but it also shows that there are many circumstances in which they would not be. What is required for extended deterrence to be relevant is that (1) concessions would increase the military capabilities of an aggressor and (2) the potential aggressor is on the verge of becoming "too powerful" (as McKay and Scott put it). But not all concessions will change the relative power of states, and even when they would, a state will have an incentive to go to war just to prevent another state from making one only if the concession would make a third state unacceptably powerful.

I pointed out in chapter 4 that the fact that bargaining does not end when war begins poses a problem for much of the cold war literature about deterrence, which assumed that it did. The analysis offered here provides a possible justification for that assumption. Paradoxically, the inability of a dissatisfied state to commit itself not to exploit the increase in its bargaining power that a concession would give it implies that a potential victim, which might have wanted to make a concession to avoid war, can commit itself not to make one. And therefore a dissatisfied state cannot hope that an attack will lead to a concession. This is an example of the more general paradox already noted more than once that the possibility of a negotiated settlement as an alternative to war can make war more likely, not less.

However, the unwillingness of a state to allow another to become more powerful cannot be inferred from a knowledge of the distribution of military capabilities among states alone. And therefore the problem of extended deterrence is to find a way to reveal that information without actually fighting.

Situations in which some states believe that another state is already on the verge of becoming unacceptably powerful prior to war are unusual, and this is a plausible way of characterizing what was distinctive about international politics during the cold war. The cold war period was different, for example, from the periods prior to both the two world wars, when Britain was concerned about two potentially expansionist states, Germany and Russia (or the Soviet Union), did not want to encourage either, but

could not oppose both simultaneously.³² Moreover, Germany did not reach the position of potential dominance that the Soviet Union achieved by helping defeat it until after a long process of expansion in Europe. And therefore, instead of joining a “balancing” coalition against Germany during the 1930s, the United States later fought a preventive war to prevent it from exploiting its control over Europe after it had defeated France.

But there is no objective way of distinguishing between distributions of power like the one prior to World War II and distributions like the one during the cold war. The only reliable indicator of when some state has crossed the threshold of tolerance of other states is the behavior of the other states, which is why the problem of extended deterrence exists.

What Next?

With the help of Clausewitz, we have seen that the explanation for the recurrence of war among the predatory rulers of early modern Europe was repeated changes in their relative military capabilities, which provided frequent opportunities for renegotiating the distribution of valuable territory but left many of them able to maintain their independence. Expectations of such changes could make it difficult to construct self-enforcing agreements among competing predators, which might, paradoxically, actually reduce the frequency of war and help preserve the independence of states. But the price was to make war more severe when it occurred.

Many of these changes in relative power were the result of competition among rulers. But even if rulers had been willing to agree to restrict their competition, changes in relative power would still have occurred as a result of domestic political and economic developments, nothing could have prevented the beneficiaries from taking advantage of them, and often the only way of resolving disagreements about their effect on the relative bargaining power of states was by fighting.³³

These wars could have been avoided had all the rulers been willing to give up their independence and subject themselves to a common ruler. But they were all too optimistic about how they would fare in future competition to agree to this, and if they had agreed, the result might have been to establish a super-predator and thus to eliminate one of the most important constraints on predatory rule in Europe.

Kant’s answer to the problem of recurring warfare among predatory

32. Since the end of the cold war, this has also been true of U.S. relations with Iraq and Iran and with Taiwan and China. For an argument that it was this feature of the cold war that Waltz tried to explain by his distinction between bipolar and multipolar systems, see Wagner 1993.

33. This is the main theme of Blainey 1988.

rulers was to eliminate the connection between rule and predation, and this answer is arguably implicit in Hobbes's limited discussion of relations among rulers as well. Structural Realists claim this would not work, but none has offered a valid argument in support of their skepticism. In the next chapter, I will reexamine this question in light of this discussion of how to explain wars among predatory rulers.