

Chapter 2: War's Inefficiency Puzzle

Chapter 0 showed why court cases are inefficient. However, we can recast that story as two countries on the verge of a military crisis. Imagine Venezuela discovers an oil deposit near its border with Colombia. Understandably, the Venezuelan government is excited; it estimates the total deposit to be worth \$80 billion. The government sends its oil company out to begin drilling.

But trouble soon arrives. Word spreads through the media of the discovery. Upon hearing the news, the Colombian government boldly declares that the oil deposit is on *its* side of the border and therefore the oil belongs to Colombia. Venezuela rejects this notion and begins drilling.

Two weeks later, political tensions reach a climax. The Colombian government mobilizes troops to the border and demands Venezuela cease all drilling operations under threat of war. In response, Venezuela sends its troops to the region. Fighting could break out at any moment.

After reviewing its military capabilities, Colombia estimates it will successfully capture the oil fields 40% of the time. However, war will kill many Colombian soldiers and damage the oil fields. After considering the price of the wasted oil and familial compensations for fallen soldiers, Colombian officials estimate its expected cost of fighting to be \$15 billion.

The Venezuelan commanders agree that Colombia will prevail 40% of the time, meaning Venezuela will win 60% of the time. Although fighting still disrupts the oil fields, Venezuela expects to lose fewer soldiers in a confrontation. Thus, Venezuela pegs its cost of fighting at \$12 billion.

On the surface, it appears the states are destined to resolve their issues on the battlefield. Colombia will win the war 40% of the time, so its expected share of the oil revenue is 40% of \$80 billion, or \$32 billion. Even after factoring in its \$15 billion in war costs, Colombia still expects a \$17 billion profit. Colombia is better off fighting than letting Venezuela have the oil.

Venezuela faces similar incentives. Since Venezuela will prevail 60% of the time, it expects to win 60% of \$80 billion in oil revenue, or \$48 billion. After subtracting its \$12 billion in costs, Venezuela expects \$36 billion in profit. Again, Venezuela prefers a war to conceding all of the oil to Colombia.

For decades, political scientists believed these calculations provided a rational explanation for war. Both sides profit from conflict. It appears insane not to fight given these circumstances.

However, upon further analysis, Venezuela and Colombia should be able to bargain their way out of war. Ownership of the oil field does not have to be an all-or-nothing affair. What if the states decided to split the oil revenue? For example, Colombia and Venezuela could set up a company that pays 60% of the revenue to Venezuela and 40% of the revenue to Colombia. If Colombia accepts the deal, it earns \$32 billion in revenue, which is \$15 billion better than had they fought. Likewise, if Venezuela accepts, it earns \$48 billion, which is \$12 billion better than the expected war outcome.

In fact, a range of bargained settlements pleases both states. As long as Colombia receives at least \$17 billion of the oil, it cannot profit from war. Similarly, if Venezuela earns at least \$36 billion, it would not want to launch a war. These minimal needs sum to just \$53 billion. Since there is \$80 billion in oil revenue to go around, the parties should reach a peaceful division without difficulty.

Where did the missing \$27 billion go? War costs ate into the revenue. It is no coincidence that Venezuela's costs (\$12 billion) and Colombia's costs (\$15 billion) sum exactly to the missing \$27 billion. These costs guarantee the existence of mutually preferable peaceful settlements. The two parties can negotiate over how that \$27 billion is divided between them. But however it is divided, even if it all goes to one side, neither side can go to war and improve its welfare.

The conflict between Venezuela and Colombia hints that bargaining could always allow states to settle conflicts short of war. We might wonder whether the result we found is indicative of a trend or a fluke convergence of the particular numbers we used in the example. To find out, we must generalize the bargaining dynamics the states face. Perhaps surprisingly, we will see that this result extends to a general framework: bargaining is always better than fighting.

The remainder of this chapter works toward proving this result. We will see three separate interpretations of the proof. To begin, we will create an algebraic formulation of the bargaining problem, which provides a clear mathematical insight: there always exists a range of settlements that leaves both sides better off than had they fought a war.

Unfortunately, the algebraic model is difficult to interpret. If you feel confused, do not despair! The next section reinterprets the problem geometrically, illustrating an example where two states bargain over where to draw a border between their two capitals. The geometric game generates a crisp visualization of the problem, which the algebraic version lacks. This will improve our understanding of these theorems as we consider more complex versions of crisis negotiations.

Finally, we will develop a game theoretical bargaining model of war. This model will become a workhorse for us in later chapters. Our attempts to explain war will ultimately attack the assumptions of this model until the peaceful result disappears. Game theory will allow us to be precise with these assumptions. Consequently, we must have expert knowledge of this model before continuing.

2.1: The Algebraic Model

Consider two states, A and B, bargaining over how to divide some good. We will let the nature of that good be ambiguous; it could be territory, money, barrels of oil, or whatever. Rather than deal with different sizes and units of the good, we standardize the good's value to 1. For example, instead of two states arguing over 16 square miles of land, they could bargain over one unit of land which just so happens to be 16 miles; the 1 effectively represents 100% of the good in its original size and in its original units. By dealing with percentages instead of specific goods, we can draw parallels between these cases.

We make a single assumption about the good: it is infinitely divisible. Thus, it is possible for one state to control .2 (or 20%) of the good while the other state controls .8 (or 80%), or for one side to have .382 and the other to have .618, and so forth.

Let p_A be state A's probability of winning in a war against B. Since p_A is a probability, it follows that $0 \leq p_A \leq 1$. (That is, p_A must be between 0% and 100%.) We will refer to p_A as state A's "power." Likewise, state B's probability of victory in a war against A (or state B's power) is p_B . Again, since p_B is a probability, we know $0 \leq p_B \leq 1$ must hold. To keep things simple, we will disallow the possibility of wars that end in partial victories or draws, though altering this assumption does not change the results. Thus, if A and B fight, one state must win and the other state must lose; mathematically, we express this as $p_A + p_B = 1$. The winner receives the entirety of the good while the loser receives nothing.

To reflect the loss of life and property destruction that war causes, state A pays a cost $c_A > 0$ and state B pays a cost $c_B > 0$ if they fight. We make no assumptions about the functional form of costs. For example, we might expect war to be a cheaper option for a state with a high probability of winning than for a state with a low probability of winning. Likewise, states that are evenly matched could expect to fight a longer war of attrition, which will ultimately cost more. Our model allows for virtually *any* relation between probability of winning and cost of fighting. The only assumption is that peace more efficiently distributes resources than war does.

Moreover, we allow the states to interpret their costs of fighting in the manner they see fit. To be explicit, c_A and c_B incorporate two facets of the conflict. First, there are the absolute costs of war. If the states fight, people die, buildings are destroyed, and the states lose out on some economic productivity. These are the physical costs of conflict.

However, the costs also account for states' *resolve*, or how much they care about the issues at stake relative to the physical costs. For example, suppose a war would result in 50,000 casualties for the United States. While Americans would not tolerate that number of lives lost to defend Botswana, they would be willing to pay that cost to defend Oregon. Thus, as a state becomes more resolved, it views its material cost of fighting as being smaller. We will discuss this concept of resolve in more depth later.

Using just the probability of victory and costs of fighting, we can calculate each state's expected utility (abbreviated EU) for war. For example, state A wins the war and takes all 1 of the good with probability p_A . With probability $1 - p_A$, state A loses and earns 0. Regardless, it pays the cost c_A . We can write this as the following equation:

$$EU_A(\text{war}) = (p_A)(1) + (1 - p_A)(0) - c_A$$

$$EU_A(\text{war}) = p_A - c_A$$

Thus, on average, state A expects to earn $p_A - c_A$ if it fights a war.

State B's expected utility for war is exactly the same, except we interchange the letter A with the letter B. That is, state B wins the war and all of the good with probability p_B . It loses the war and receives 0 with probability $1 - p_B$. Either way, it pays the cost c_B . As an equation:

$$EU_B(\text{war}) = (p_B)(1) + (1 - p_B)(0) - c_B$$

$$EU_B(\text{war}) = p_B - c_B$$

Given these assumptions, do any negotiated settlements provide a viable alternative to war for both sides simultaneously? Let x represent state A's share of a possible settlement. Recalling back to the standardization of the good, x represent the percentage of the good state A earns. State A cannot improve its outcome by declaring war if its share of the bargained resolution is greater than or equal to its expected utility for fighting. Thus, A accepts any resolution x that meets the following condition:

$$x \geq p_A - c_A$$

Likewise, B is at least as well off as if it had fought a war if its share of the bargained resolution is greater than or equal to its expected utility for fighting. Since the good is worth 1 and B receives everything that A did not take, its share of a possible settlement is $1 - x$. Thus, B accepts any remainder of the division $1 - x$ that meets the following condition:

$$1 - x \geq p_B - c_B$$

To keep everything in terms of just x , we can rearrange that expression as follows:

$$1 - x \geq p_B - c_B$$

$$x \leq 1 - p_B + c_B$$

Since x is A's share of the bargain, the rearranged expression has a natural interpretation: B would rather fight than allow A to take more than $1 - p_B + c_B$.

Combining the acceptable offer inequalities from state A and state B, we know there are viable alternatives to war if there exists an x that meets the following requirements:

$$x \geq p_A - c_A$$

$$x \leq 1 - p_B + c_B$$

$$p_A - c_A \leq x \leq 1 - p_B + c_B$$

Thus, as long as $p_A - c_A \leq 1 - p_B + c_B$, such an x is guaranteed to exist. Although we may appear to be stuck here, our assumptions give us one more trick to use. Recall that war must result in state A or state B winning. Put formally:

$$p_A + p_B = 1$$

$$p_B = 1 - p_A$$

In words, the probability B wins the war is 1 minus the probability A wins the war. Having solved for p_B in this manner, we can substitute $1 - p_A$ into the previous inequality:

$$p_A - c_A \leq 1 - p_B + c_B$$

$$p_B = 1 - p_A$$

$$p_A - c_A \leq 1 - (1 - p_A) + c_B$$

$$p_A - c_A \leq 1 - 1 + p_A + c_B$$

$$-c_A \leq c_B$$

$$c_A + c_B \geq 0$$

So a bargained resolution must exist if sum of c_A and c_B is greater than or equal to 0. But recall that both c_A and c_B are individually greater than zero. Thus, if we sum them together, we end up with a number greater than 0. We can write that as follows:

$$c_A + c_B > 0$$

Therefore, we know $c_A + c_B \geq 0$ must hold. In turn, a bargained resolution must exist!

Put differently, the oil example from earlier was no fluke; there always exists a range of peaceful settlements that leave the sides at least as well off as if they had fought a war. The settlement x must be at least $p_A - c_A$ but no more than $1 - p_B + c_B$, and we know the states can always locate such an x because of the positive costs of war.

2.2: The Geometric Model

The algebraic model provided an interesting result: peace is mutually preferable to war. However, it is hard to interpret those results. The proof ended with $c_A + c_B > 0$; such mathematical statements provide little intuitive understanding of why states ought to bargain.

Thus, in this section, we turn to a geometric interpretation of our results. Essentially, we will morph the algebraic statements into geometric pictures. The visualization helps explain why the states ought to settle rather than fight.

Let's start by thinking of possible values for x , the proposed division of the good, as a number line. Since x must be between 0 and 1, the line should cover that distance:

0

1

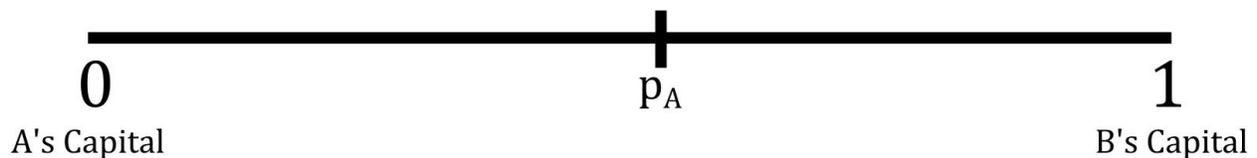
Think of this line as a strip of land. $x = 0$ represents state A's capital; $x = 1$ represents state B's capital. Each state wants as much of the land as it can take. Thus, the closer the states draw the border to 1, the happier A is. On the other hand, state B wants to place the border as close to 0 as possible.

We can label the capitals accordingly:



In sum, A wants to conquer land closer to B's capital while B wants to conquer land closer to A's capital. Keep in mind, however, that this model also applies to other types of bargaining objects. Indeed, in later chapters we will discuss bargaining situations between two countries that do not even border each other. A's capital merely reflects A's least preferred outcome (and consequently B's most preferred outcome), while B's capital represents A's most preferred outcome (and B's least preferred outcome).

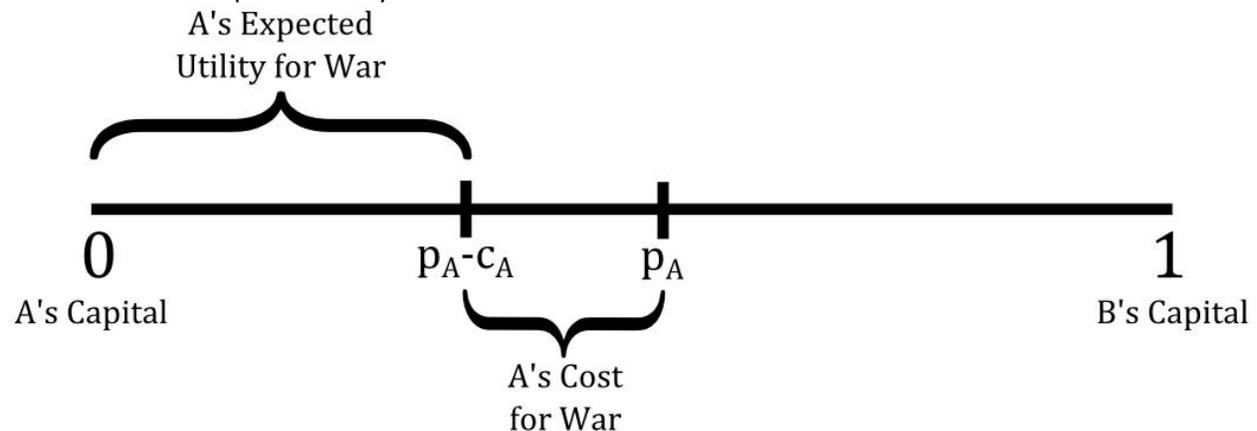
For now, though, we will stick to territory. Let's think about the types of borders the states would prefer to war. If the states fight a war, A wins with probability p_A and will draw the border $x = 1$. With probability $1 - p_A$, B wins the war and chooses a border of $x = 0$. Consequently, in expectation, war produces a border of $x = p_A$:



The strip of land to the left of p_A represents A's expected share of the territory. Here, that amount equals p_A . The strip of land to the right corresponds to B's expected share. Since B earns everything between p_A and 1, that amount is $1 - p_A$. Note that the drawn location of p_A is generalized; although it appears to be slightly further than half way, it could actually be anywhere on the line.

Before factoring in the costs of war, it is clear that A would be happy to divide the territory at any point to the right of p_A , since war would draw the border at p_A in expectation. Likewise, B would be happy to divide the territory at any point to the left of p_A , since that pushes the border further from B's capital than war does.

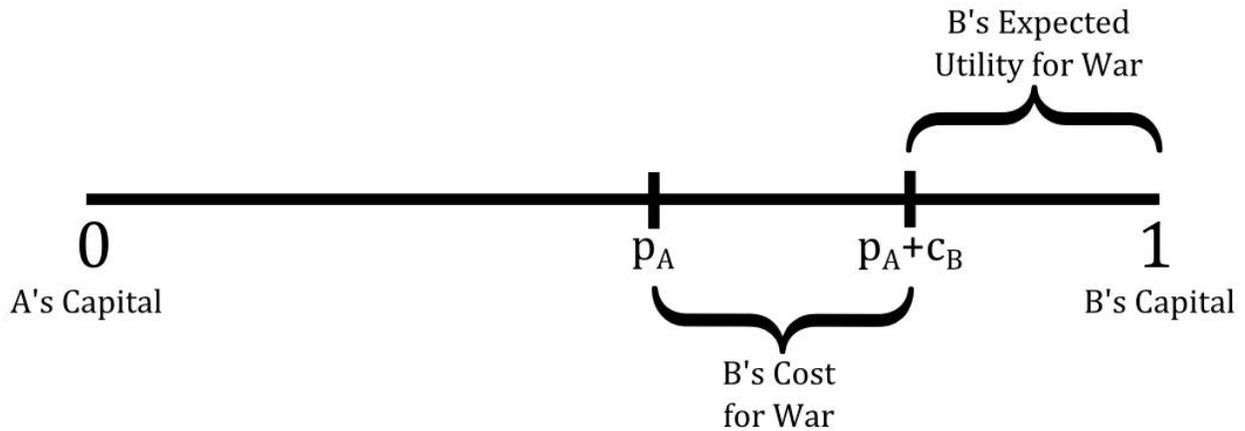
However, war is a costly option for both states. If they fight, A earns an expected territorial share of p_A but must pay a cost of fighting c_A . Thus, its expected utility for war is not p_A , but rather $p_A - c_A$. We can illustrate A's expected utility as follows:



Obviously, A is still pleased to draw the border to the right of p_A . But these costs also mean A prefers a border in between $p_A - c_A$ and p_A to fighting a war. Although war ultimately produces a border closer to A's ideal outcome than $p_A - c_A$, the costs of fighting make conflict not worth the expense. Thus, all told, A is willing to accept any settlement that draws the border to the right of $p_A - c_A$.

B's preferences are similar. War is also a costly option for B. If the states fight, B earns a territorial share of $1 - p_A$ in expectation but still pays the cost c_B . Thus, the costs of war push B's expected utility closer to B's capital.

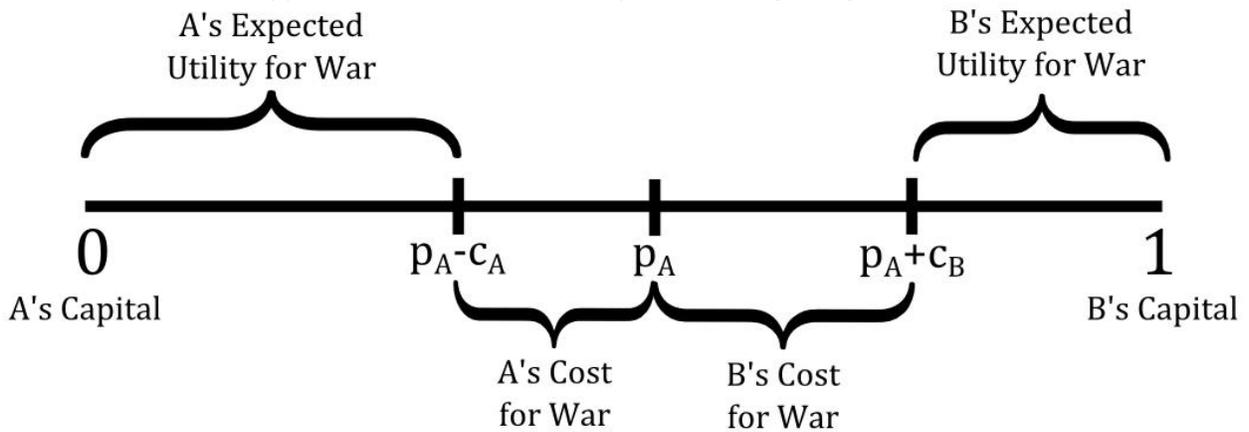
We can illustrate B's preferences like this:



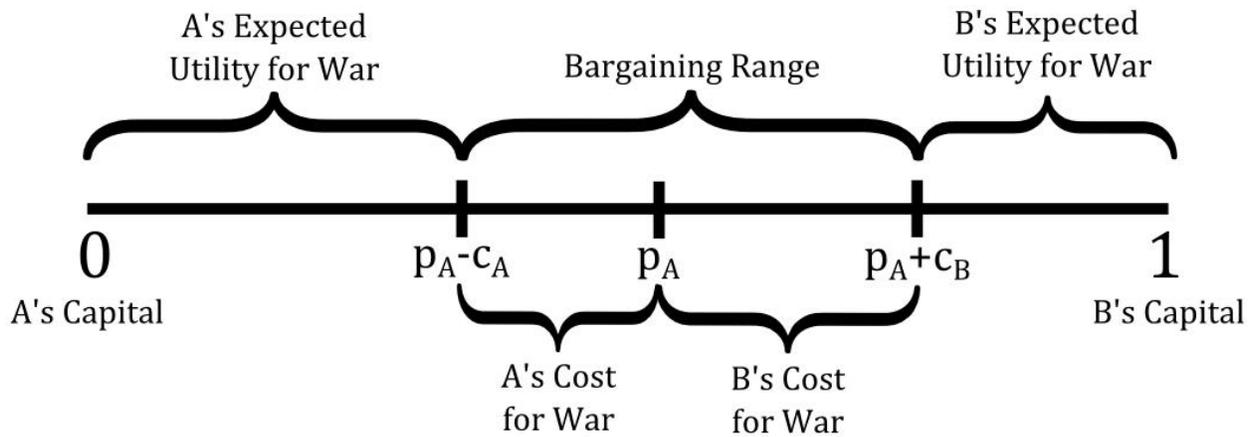
This time, any border to the left of $p_A + c_B$ satisfies B.

The plus sign in front of c_B might be counterintuitive. Despite costs being bad for B, we must add c_B to p_A to draw B's effective outcome closer to its capital and further away from its ideal outcome. Since B's expected utility for war is the space in between 1 and $p_A + c_B$, its expected utility equals $1 - (p_A + c_B)$, or $1 - p_A - c_B$. Thus, even though war produces an expected border at p_A , B is still willing to accept borders drawn between p_A and $p_A + c_B$.

But notice what happens when we combine the previous images together:



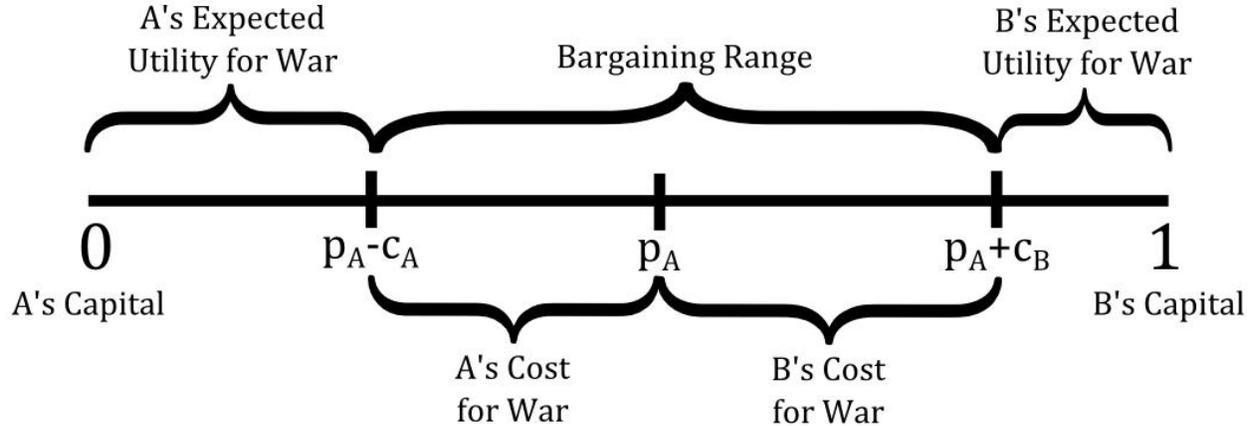
To satisfy A, B must draw the border to the right of $p_A - c_A$; to satisfy B, A must draw the border to the left of $p_A + c_B$. Thus, any border between $p_A - c_A$ and $p_A + c_B$ satisfies both parties. We call this the *bargaining range*:



This directly corresponds to what we saw in the algebraic version of the model. Recall that a viable alternative to war was any compromise x that satisfied the requirement $p_A - c_A \leq x \leq 1 - p_B + c_B$. The geometric model simply shows us what such an x means; the bargaining range is all of the values for x that fulfill those requirements.

The geometric interpretation also allows us to better understand how a state's resolve corresponds to its cost. Suppose the above example involved two countries fighting over valuable territory; perhaps the space between them contains some natural resource like oil. Consequently, they are willing to pay great costs to take control of the land.

Alternatively, suppose these same states were looking at a different strip of territory between their capitals. This time, the land is arid and not particularly useful. Although the states have the same capabilities and will endure the same absolute costs of fighting, they will be less resolved over the issue since the land is relatively worthless. As such, the relative costs of fighting will be greater:



Note that the size of the bargaining range remains exactly $c_A + c_B$. Thus, as states become less resolved over the issues, they are willing to agree on more bargained resolutions. The expanded size of the bargaining range reflects this.

The same is true in terms of absolute costs. Consider a border dispute. In the first case, both states have weak military forces. Consequently, they cannot inflict much damage to each other. In the second case, both states have strong militaries and have nuclear capabilities. War is a much costlier option for both parties here. As such, the bargaining range is much larger in the second case than the first even though the states are fighting over the same piece of territory in both cases.

Finally, we can also incorporate hawkish and dovish preferences into these cost functions. Hawkish states do not find killing people to be as morally reprehensible as dovish states do. In turn, states with dovish cultures face a higher perceived cost of war than hawkish states. Unfortunately, this leaves dovish states in a vulnerable position, as hawkish states can take advantage of their reluctance to fight.

As such, dovish states may want to act as hawkish states to protect their share of the bargain. Chapter 4 will investigate whether bluffing in this manner can cause war.

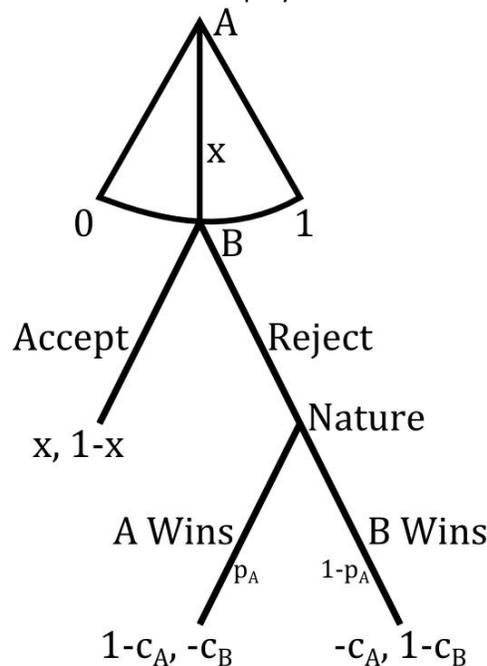
Although the geometric approach to bargaining provides us with a clear conceptual framework, we lose out on a bit of precision. As just mentioned, we will consider modifications to the bargaining situation. Perhaps states may be uncertain of each other's capabilities or resolve. Power could shift over the course of time. States may only be able to implement particular divisions of the good. Unfortunately, the algebraic and geometric versions of the model cannot adequately describe such rich environments. As such, we must turn to a game theoretical approach.

2.3: The Game Theoretical Model

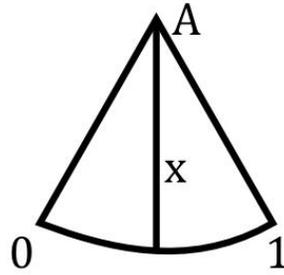
Transitioning to game theory allows us to take advantage of the tools game theorists have been developing for decades. The one downside is that we must impose slightly more structure to the interaction to work in a game theoretical world. Rather than searching for divisions that satisfy both parties, we will suppose state A is a *status quo* state; it owns the entire good, which we still standardize as worth 1. Meanwhile, state B covets the good and is potentially willing to fight a war if state A does not concede enough of it.

More precisely, the interaction is as follows. State A begins the game by offering state B a division of the good. As before, we will call the amount A keeps x . B observes A's demand and accepts it or rejects it. If B accepts, the states settle the conflict peacefully. If B rejects, the states fight a war in which A prevails with probability p_A and B prevails with probability p_B .

We can use a game tree to illustrate the flow of play:

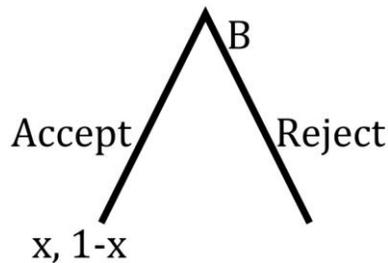


Since most of our future chapters utilize game trees like this one, we ought to spend a moment understanding what everything means. Let's start at the top:



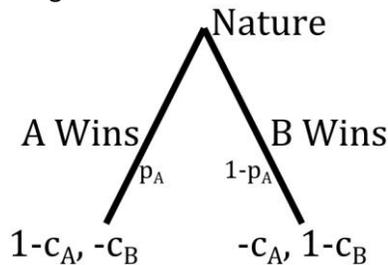
State A starts by making an offer x . The curved line indicates that A chooses an amount between 0 and 1. Thus, A is free to pick *any* value for x that satisfies those constraints, whether it be 0, .1, .244, .76, 1, or whatever.

Following that, B makes its move:



Here, B has two choices. If B accepts, the states receive the payoffs listed. By convention, state A receives the first number and state B receives the second. Thus, A receives x and B receives $1 - x$.

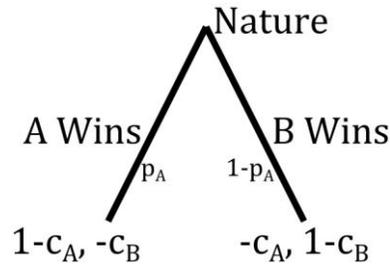
If B rejects, we move to the final stage:



Nature acts as a computerized randomizer. With probability p_A , it selects A as the winner of the war. As the victor, A can impose any settlement it wishes. Since A wants to maximize its own share of the territory, it assigns the entire strip of the territory (worth 1) to itself. But A still pays the cost of war, leaving it with an overall payoff of $1 - c_A$. B, meanwhile, receives none of the territory but pays the cost to fight, giving it payoff of just $-c_B$. With probability $1 - p_A$, B wins the war, and the same logic applies in reverse.

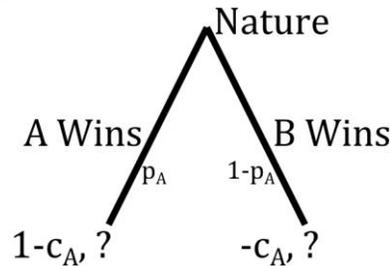
How do we solve this game? There may be temptation to start at the top and work downward. After all, the states move in that order. It stands to reason we should solve it that way as well. However, the optimal move at the beginning depends on how today's actions affect tomorrow's behavior. A state cannot know what is optimal at the beginning unless it anticipates how the rest of the interaction will play out. Thus, we must start at the *end* and work our way backward. Game theorists call this solution concept *backward induction*. Although we will not fully explore backward induction's power in this book, we can nevertheless apply it to this model. (For more on backward induction, consider reading [Game Theory 101: The Complete Textbook](#).)

Fortunately, the process of solving the game is fairly painless. To start, recall that the interaction ends with nature randomly choosing whether A or B wins:



Although the states do not know who will actually prevail in the conflict, they can calculate their expected utilities for fighting. To do this, as we have done before, we simply sum each actor's possible payoffs multiplied by the probability each outcome actually occurs.

Let's start with state A's payoffs:



With probability p_A , A wins and earns $1 - c_A$. With probability $1 - p_A$, A loses the war and earns $-c_A$. To calculate A's expected utility for war, we multiply these probabilities by their associated payoffs and sum them together. Thus, state A's expected utility for war equals:

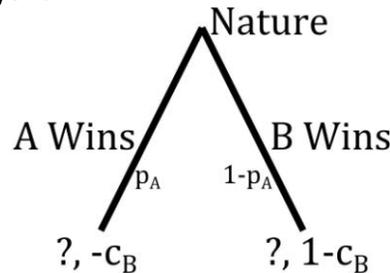
$$EU_A(\text{war}) = (p_A)(1 - c_A) + (1 - p_A)(-c_A)$$

$$EU_A(\text{war}) = p_A - p_A c_A - c_A + p_A c_A$$

$$EU_A(\text{war}) = p_A - c_A$$

Note that this is exactly the same war payoff A had in the algebraic version of the model. The benefit of the game tree is that we see that A never actually earns a payoff of $p_A - c_A$ at the end of a war if it fights. Instead, $p_A - c_A$ reflects state A's expectation for nature's move. Sometimes, nature is friendly, allows A to win, and thereby gives A more. Sometimes, nature is less friendly, forces A to lose, and thereby gives A much less. But the weighted average of these two outcomes is $p_A - c_A$.

Now let's switch to state B's payoffs:



Here, state B loses and earns $-c_B$ with probability p_A , while it wins and earns $1 - c_B$ with probability $1 - p_A$. As an equation:

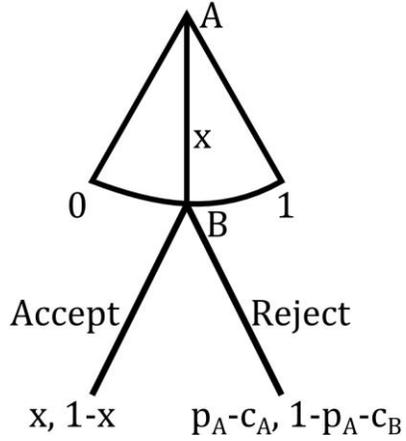
$$EU_B(\text{war}) = (p_A)(-c_B) + (1 - p_A)(1 - c_B)$$

$$EU_B(\text{war}) = -p_A c_B + 1 - c_B - p_A + p_A c_B$$

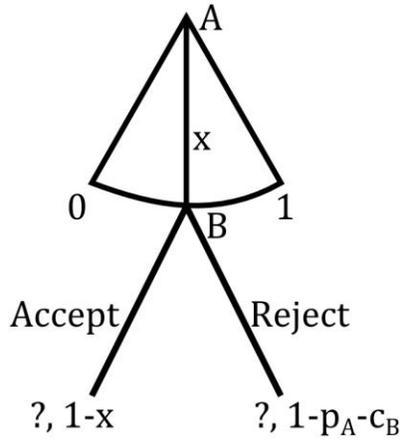
$$EU_B(\text{war}) = 1 - p_A - c_B$$

Thus, B earns $1 - p_A - c_B$ in expectation if it rejects A's offer and fights a war.

Now that we have both states' expected utilities for war, we can erase nature's move and make these payoffs the ultimate outcome for B rejecting:



With this reduced game, we can now see which types of offers B is willing to accept. Let's focus on B's payoffs:



B can accept any offer $1 - x$ that is at least as good as $1 - p_A - c_B$, its expected utility for war. As an inequality:

$$1 - x \geq 1 - p_A - c_B$$

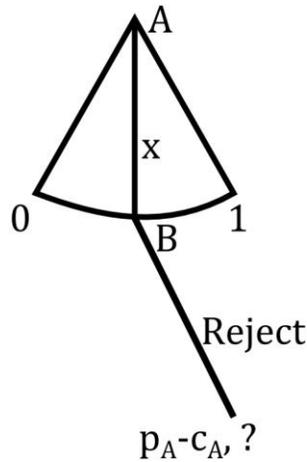
$$-x \geq -p_A - c_B$$

$$x \leq p_A + c_B$$

Thus, B is willing to accept x as long as it is less than or equal to $p_A + c_B$. That is, if A demands more than $p_A + c_B$, B must reject it.

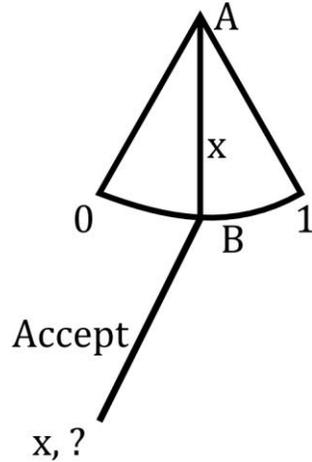
Finally, we move back to state A's decision. State A has infinitely many values to choose from: 0, .1, 1/3, .5, .666662, .91, and so forth. Yet, ultimately, these values fall into one of two categories: demands acceptable to B and demands unacceptable to B.

Suppose A selects an x greater than $p_A + c_B$. Then B rejects. Using the game tree, we can locate state A's payoff for such a scenario:



Consequently, we can bundle all of these scenarios into one expected utility. If state A makes an unacceptable offer to B—whether it is slightly unacceptable or extremely unacceptable—B always fights a war, and A winds up with $p_A - c_A$.

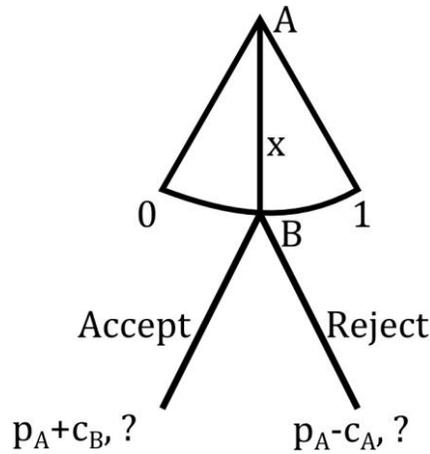
In contrast, suppose A demanded $x \leq p_A + c_B$. Now state B accepts. Here is that outcome:



This time, A simply earns x , which is the size of its peaceful demand. This variable payoff complicates matters. When B rejected, A earned the same payoff every time. Here, however, A's payoff is different for every acceptable offer it makes.

So which is A's best acceptable offer? Note that A wants to keep as much of the good as it can. Thus, if A prefers inducing B to accept its demand, A wants that demand to be as beneficial to itself as possible. Since B accepts any $x \leq p_A + c_B$, the largest x that B is willing to accept is $x = p_A + c_B$. In turn, if A ultimately wants to make an acceptable demand, the best acceptable demand it can make is $x = p_A + c_B$; any smaller value for x needlessly gives more of the good to B.

Although we started with an infinite number of possible optimal demands (all of which were between 0 and 1), we have narrowed A's demand to $x = p_A + c_B$ or any $x > p_A + c_B$. Since we know the best acceptable demand A can make is $x = p_A + c_B$, let's insert that substitution into the game tree. And because state A controls the offer, let's also isolate A's payoffs:



Thus, A should make the acceptable offer $x = p_A + c_B$ if its expected utility for doing so is at least as great as A's expected utility for inducing B to reject. As an inequality:

$$p_A + c_B \geq p_A - c_A$$

$$c_B \geq -c_A$$

$$c_A + c_B \geq 0$$

But as we saw in an earlier section, we know this inequality must hold because both c_A and c_B are greater than 0 by definition.

Therefore, in the outcome of the game, A demands $x = p_A + c_B$ and leaves $1 - p_A - c_B$ for B. B accepts the offer, and the states avoid war once again.

We call this outcome the *equilibrium* of the game. Although the payoffs might not be balanced in the way the "equilibrium" might imply, we use that word because such a set of strategies is stable. Neither side can change what they were planning on doing and expect to earn a greater average payoff.

The concept of equilibrium is compelling. After all, if states are intelligent, they ought to be maximizing the quality of their outcomes. Finding equilibria ensures that each actor is doing the best it possibly can given that another actor is attempting to do the same. We will be working extensively with this concept in upcoming chapters.

That aside, it is worth comparing the specific result in the game theoretical model to the more general results in the algebraic and geometric models. The first two models predicted the resolution would be some agreement at least as great as $p_A - c_A$ but no greater than $p_A + c_B$. In contrast, the game theoretical model specifically guesses that $x = p_A + c_B$ will be the result. What accounts for the difference?

Note that the game theoretical model makes an important assumption the others do not: state A chooses its demand. We justified this by assuming that A controls all of the good to start with. Thus, when B initiates negotiations, A can choose exactly how much to leave on the table for B to accept or reject. Since A wants to keep as much for itself as possible, it selects the exact amount that will satisfy B. Although B earns less than it would have had A been more generous with its offer, B cannot improve its outcome by fighting. Giving A control of the demands allowed A to reach the point of the bargaining range most advantageous to it. It should not be at all surprising that A takes as much as B is willing to let it.

2.4: What Is the Puzzle?

In each of the models, we saw that practical alternatives to war always exist. As such, if states reach an impasse in bargaining, it cannot be because no settlement is mutually preferable to war. Instead, it

must be that states fail to recognize these settlements or refuse to believe they can be implemented in an effective manner.

The existence of such deals immediately cast doubt on the popular explanations for many wars. For example, consider the 2011 Libyan Civil War. Conventional wisdom says that the war started because of Muammar Gaddafi's oppression of his citizens and massive inequality within the country. While these grievances certainly existed, they do *not* explain why the war broke out. After all, Gaddafi's regime could have simply relaxed the level of oppression and offered economic concessions to appease the opposition.

Similarly, the standard explanation for the Persian Gulf War is that Saddam Hussein invaded Kuwait and the United States would not tolerate such aggression. Again, though, this does not explain why war occurred. Indeed, Saddam could have simply stolen a handful of oil fields from Kuwait instead launching a full-on invasion. While this would have undoubtedly upset Kuwait, the United States, and most of the rest of the world, it is questionable whether tensions would have escalated as far as they did if Saddam had acted less aggressively.

Overall, popular explanations for war generally point to some grievance between the two fighting parties. This is useful to some degree. Grievances are certainly *necessary* for war—if no disagreement exists, no reason to fight exists—but they are not *sufficient* for war. Grievances exist all over. Why, then, does war break out over some grievances but not others?

War's inefficiency puzzle therefore asks why states sometimes choose to resolve their differences with inefficient fighting when they could simply select one of these peaceful and mutually preferable alternatives. That is, we are seeking explicit reasons why states cannot locate one of these peaceful settlements or cannot effectively implement them.

In turn, a *rationalist explanation for war* answers war's inefficiency puzzle while still assuming the states only want to maximize their share of the goods at stake minus potential costs of fighting. Over the course of this chapter, we made some strong assumptions about the states' knowledge of each other and the structure of power over time. If we weaken these assumptions, the states may rationally end up fighting each other. The next few chapters explore four of these explanations: preventive war, private information and incentives to misrepresent, issue indivisibility, and preemptive war.

In the baseline model, we looked at a snapshot in time, during which power stayed static; state A always won the war with probability p_A and state B always won with probability $1 - p_A$. However, relative military power fluctuates over the years. A weak country today can develop its economic base, produce more tanks, begin research into nuclear weapons, and become more threatening to its rivals in the future. Thus, declining states might want to quash rising states before the latter becomes a problem. Political scientists call this *preventive war (or preventative war)*, and we cover it in the next chapter.

Moving on, the states were perfectly aware of each other's military capabilities and resolve in the baseline model. This is a strong assumption. In reality, military commanders have private information about their armies' strengths and weaknesses. Perhaps the lack of knowledge causes states to overestimate the attractiveness of war, which in turn leads to fighting. Chapter four explores such a scenario and shows how the possibility of bluffing sabotages the bargaining process.

The fifth chapter relaxes the infinitely divisible nature of the good the states bargain over. Although states can divvy up land, money, and natural resources with ease, other issues may not have natural divisions. For example, states cannot effectively split sovereignty of a country. Either John can be king or Mark can, but they both cannot simultaneously be *the* king. Political scientists call this restriction *issue indivisibility*. But whether issue indivisibilities actually exist is still a matter of debate. We will tackle it in chapter five.

The baseline model also assumes that power remains static regardless of which state starts the war. If A initiates, it wins with probability p_A ; if B attacks first, A still wins with probability p_A . However, first strike advantages might exist. After all, the initiator may benefit from surprising the other party and

dictating when and where the states fight battles. If these advantages are too great, the temptation to defect from a settlement will keep states from ever sitting down at the bargaining table. Political scientists call this *preemptive war*, and we cover it in the sixth chapter.

2.5: Further Reading

This chapter diagrammed the fundamental puzzle of war that James Fearon presented in “Rationalist Explanations for War.” You can view the original article [here](#).