Chapter 2: War's Inefficiency Puzzle

Imagine Venezuela discovered a new oil deposit near the border of Colombia. Understandably, the Venezuelan government is excited; it estimates the total deposit to be worth $80 billion.

But trouble soon arrives. 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 to 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 of Colombian soldiers and damage the oil fields. Colombian officials estimate the 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. If Colombia wins the war 40% of the time, its expected share of the oil revenue is $80 billion times 40%, or $32 billion. Even after factoring in Colombia's $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. If Venezuela prevails 60% of the time, it expects to win $80 billion times 60% of 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. If both sides expect a net profit from war, conflict seems inevitable.

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, in earns $48 billion, or $12 billion better than the expected war outcome.

In fact, a range of bargained settlements appeases 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. Since there is $80 billion in oil revenue to go around and these war payoffs sum to $53 billion, the parties should reach a peaceful division without difficult.

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 the peaceful settlements.

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 for this specific example was a fluke for the particular numbers we chose or if it is indicative of a trend. To find out, we must generalize the bargaining dynamic the states face. Perhaps surprisingly, we will see that this result extends to a general framework.

The remainder of this chapter works toward proving this result. I offer three separate interpretations of the proof. We will begin with 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.

Next, we will consider 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 will help our understanding as we engage in more complex versions of the problem.

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 it breaks. 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 should standardize the good's value to 1. Thus, instead of two states arguing over 16 square miles of land, they instead bargain over one unit of land which just so happens to be 16 miles. In effect, the 1 represents 100% of the good in its original size and in its original unit. By dealing with percentages instead of specific goods, we can draw parallels among all of these cases.

We make but 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 $0 \leq p_A \leq 1$. 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. For now, we will disallow the possibility of wars that in 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.

However, 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 also 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 war is not free.

Moreover, we allow the states to interpret its costs of fighting in the manner it sees fit. More explicitly, 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.

However, the costs also factor in states *resolve*, or how much they care about the issues at stake. For example, suppose a war would result in 50,000 causalities 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 absolute cost of fighting as being smaller. We will discuss this concept 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 this as the following equation:

$EU_A(war) = (p_A)(1) + (1 – p_A)(0) – c_A$

$EU_A(war) = p_A - c_A$

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:

$EU_B(war) = (p_B)(1) + (1 - p_B)(0) - c_B$
$EU_B(war) = p_B - c_B$

Given these assumptions, do any negotiated settlements leave both sides better off than had they fought a war? 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 is at least as well off than if it had fought a 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 than 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 resolution $1 - x$ 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:

$p_A + p_B = 1$
$p_B = 1 - p_A$

Having solved for p_B, we can substitute $1 - p_A$ into the previous inequality:

$p_A - c_A \leq 1 - p_B + c_B$
$p_B = 1 - p_A$
$p_A - c_A \leq 1 - (1 - p_A) + c_B$
$p_A - c_A \leq 1 - 1 + p_A + c_B$
$-c_A \leq c_B$
$0 \leq c_A + c_B$

So a bargained resolution must exist if 0 is less than or equal to the sum of c_A and c_B. 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:

$c_A + c_B > 0$

Thus, we know $c_A + c_B > 0$ 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 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.

2.2: The Geometric Model

The algebraic model provided an interesting result: there exists a range of settlements that provide each side with at least its payoff for 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 turn 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 as a number line. Since x must be between 0 and 1, our line should cover that distance:

Think of this strip of land as territory. x = 0 represents state A's capital; x = 1 represents state B's capital. The states want to maximize their share of the land. Thus, the closer the states draw the border to 1, the happier A is; after all, x = 1 is A's ideal point. On the other hand, state B wants to place the border as close to 0 as possible.

Recall that A wants to maximize its share of this number line; if x = 1, A receives all of the good. In contrast, B wants to minimize A's share of the number line; if x = 0, A receives none of the good, which means B receives all of it. Thus, we can refer to x = 1 as A's *ideal point* while x = 0 is B's ideal point:

Essentially, A wants to conquer lands closer to B's capital while B wants to conquer lands closer to A's capital.

Now 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 $x = p_A$:

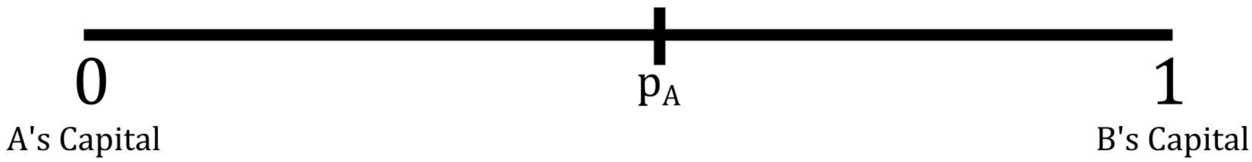

The strip of land to the left of p_A represents A's expected share of the territory. Here, that amount equals p_A. The strip of land to the right corresponds to B's expected share. Since B earns everything between p_A and 1, that amount is $1 – p_A$.

However, war is a costly option for A. If the states 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:

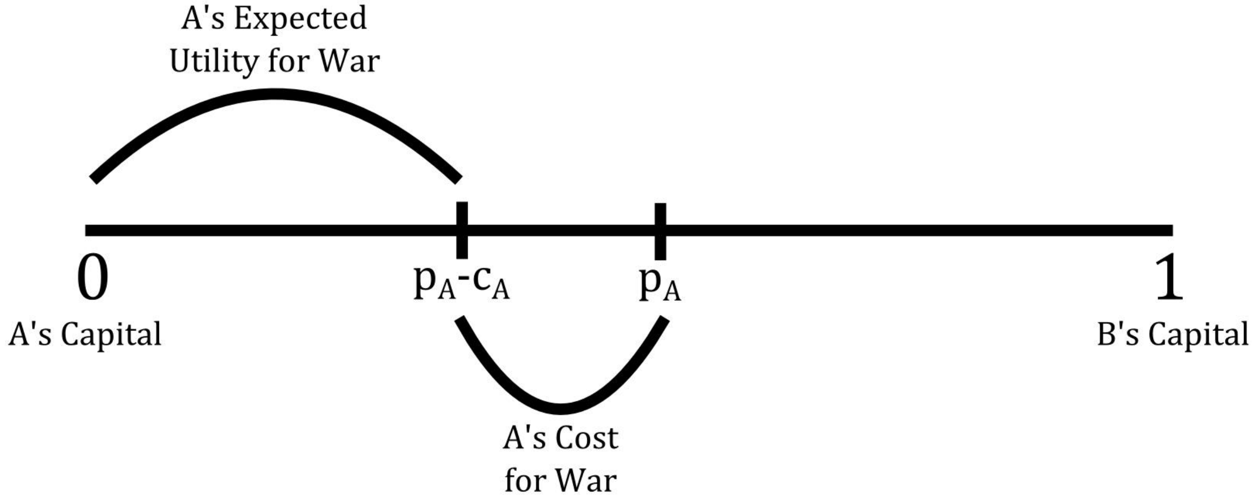

Obviously, A is happy to draw the border to the right of p_A. But these costs also mean A prefers a border in between $p_A – c_A$ and p_A to fighting a war. Although war ultimately produces a border of p_A which is close to A's ideal point than $p_A – c_A$, the costs of fighting make war not worth the expense. Thus, A is willing to accept any settlement that draws the border to the right of $p_A – c_A$.

We can draw B's preferences in a similar way. War is also a costly option for B. If the states fight, it earns a territorial share of $1 – p_A$ in expectation. However, B must pay a cost c_B to fight. Thus, the costs of war push B's expected utility closer to B's capital. We can see it like this:

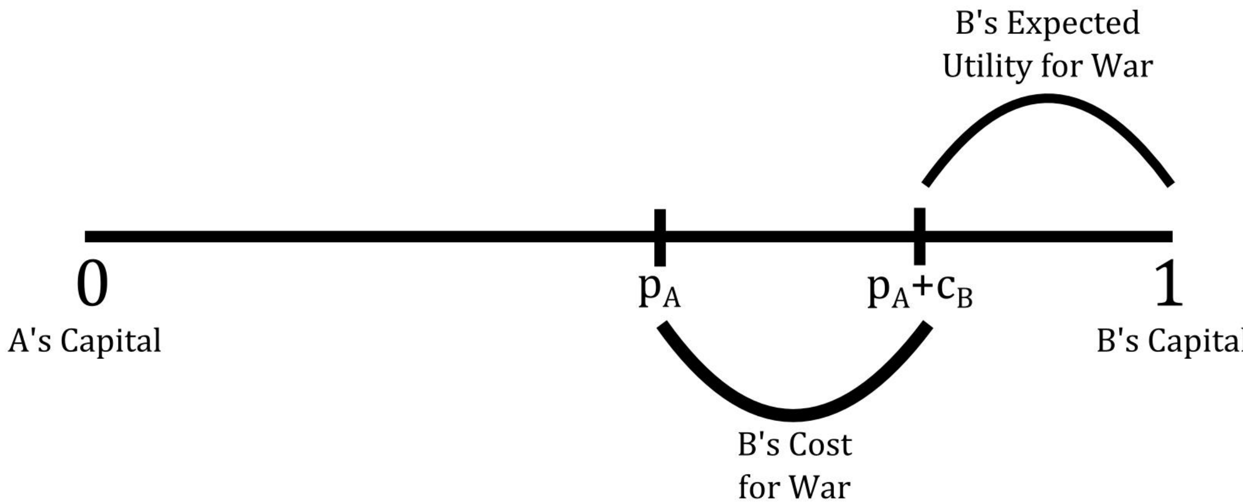

This time, any border to the left of $p_A + c_B$ satisfies. Despite costs being bad for B, we must add c_B to p_A to draw B's effective outcome closer to its capital and further away from its ideal point. Since B's expected utility for war is the space in between 1 and $p_A + c_B$, its expected utility 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 draw in between p_A and $p_A + c_B$.

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

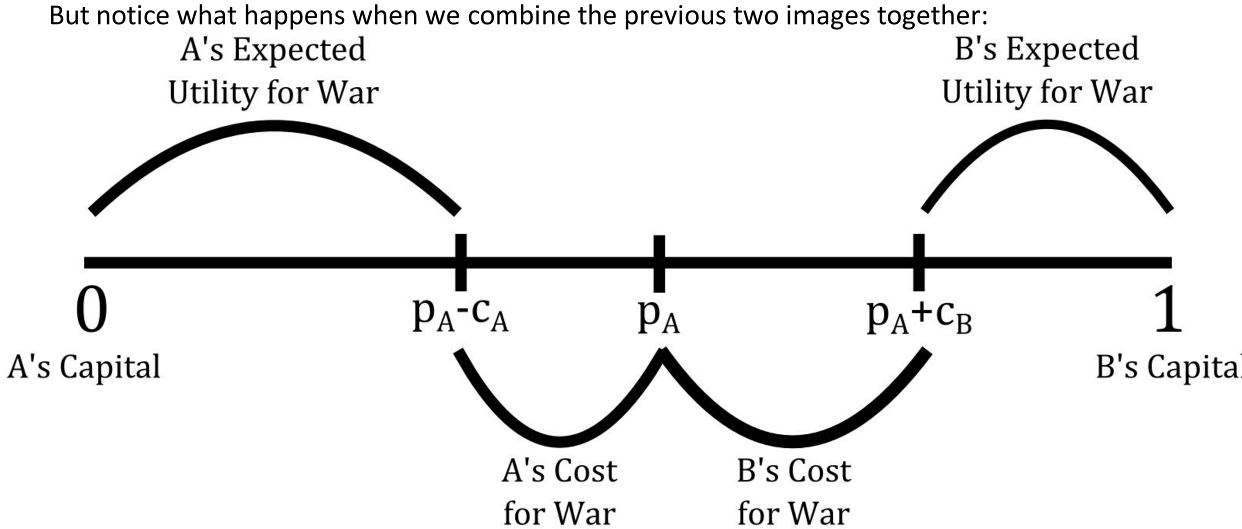

To appease A, B must draw the border to the right of $p_A - c_A$; to appease B, A must draw the border to the left of $p_A + c_B$. Thus, any border between $p_A - c_A$ and $p_A + c_B$ satisfies both parties. We call this the *bargaining range*:

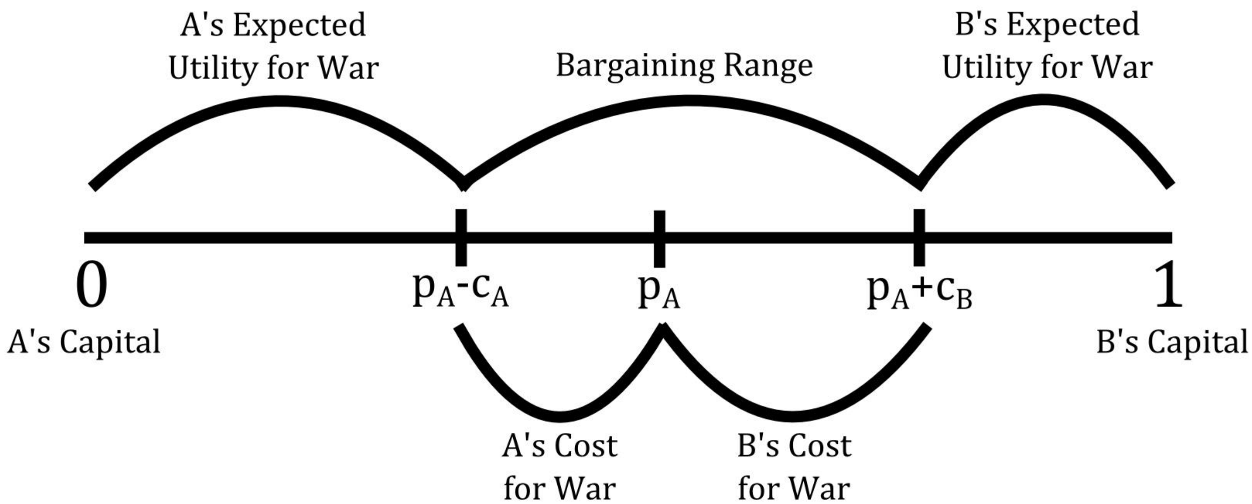

This directly corresponds to what we saw in the algebraic version of the model. Recall that a viable alternative to war was any compromise x that satisfied the requirement $p_A - c_A \leq x \leq 1 - p_B + c_B$. The geometric model simply shows us what such an x means; the bargaining range is all of the 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, however, the land is rugged 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 is relatively worthless. As such, the relative costs of fighting will be greater:

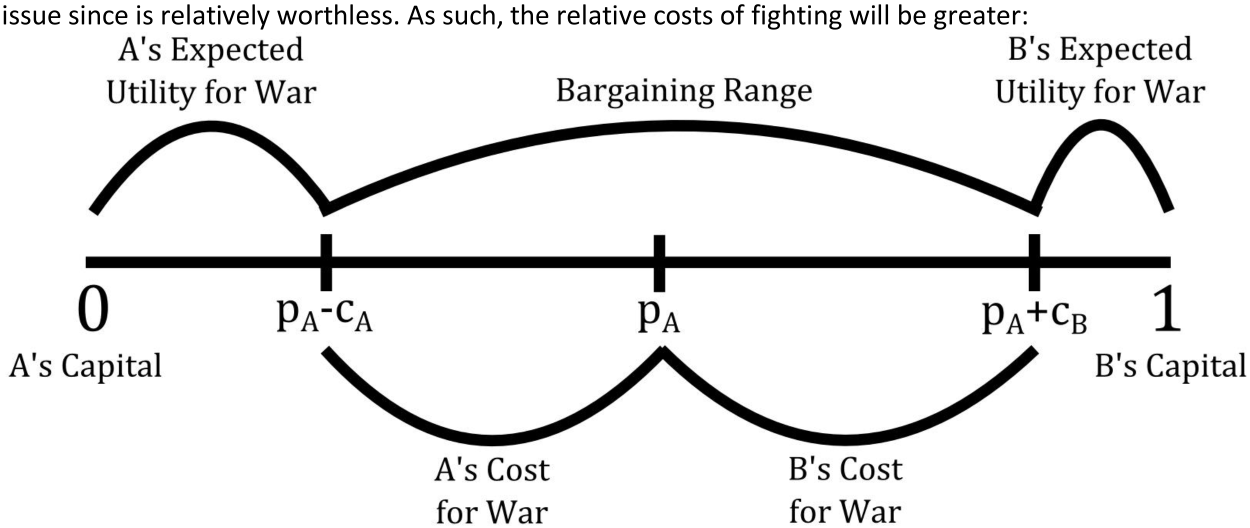

Thus, as states become less resolved over the issues, they are willing to agree on more bargained resolutions. The expanded size of the bargaining range reflects this.

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.

Although the geometric approach to bargaining provides us with a clear conceptual framework, we lose out on a bit of precision. In future chapters, we will consider modifications to the bargaining situation; states may be uncertain of each other's capabilities or resolve, power could shift over the course of time, or 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

Turning 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 owns the entire good, which is still worth 1. Meanwhile, state B covets the good and is willing to fight a war if state A does not appease 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 offer 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:

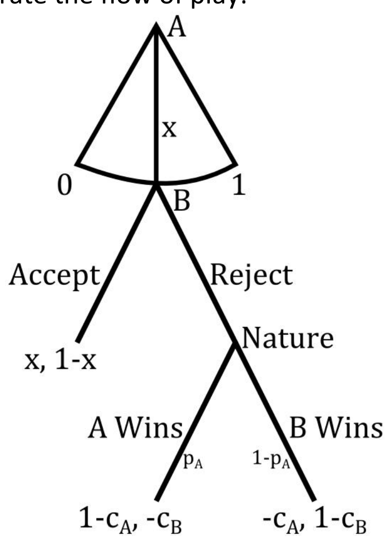

Since the most of our future chapters utilize game trees like this one, we ought to spend a moment understanding what everything means. Let's start at the top:

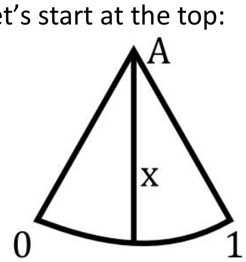

State A starts by making an offer x; the curved line between 0 and 1 indicates that that x must no smaller than 0 but no greater than 1.

Following that, B makes its move:

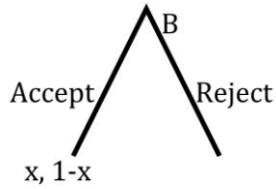

Here, B has two choices. If B accepts, the states receive the payoffs listed. By convention, state A receives the first number and state B receives the second. If B rejects, we move to the final stage:

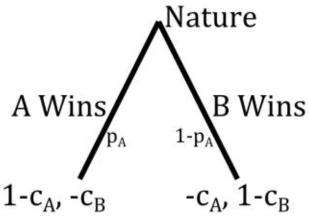

Nature acts as a computerized randomizer. With probability p_A, it selects A as the winner of the war and assigns payoffs of $1 - c_A$ to A and of $-c_B$ to B. With probability $1 - p_A$, it picks B as the winner and gives payoffs of $-c_A$ to A and $1 - c_B$ to B.

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, for the states to make smart moves at the beginning, they need to anticipate how the other state will respond to particular actions.

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 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:

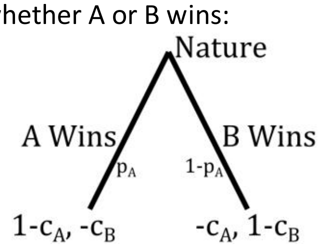

Although the states do not know who will actually prevail in the conflict, they can calculate their expected utilities for fighting. To do this, as we have done before, we simply sum each actor's possible payoffs multiplied by the probability each outcome actually occurs.

Let's start with state A's payoffs:

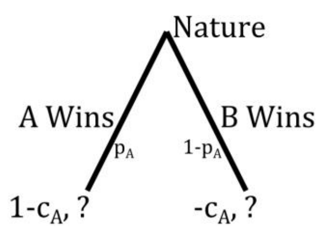

With probability p_A, A wins and earns $1 - c_A$. With probability $1 - p_A$, A loses the war and earns $-c_A$. Thus, state A's expected utility for war equals:

$EU_A(war) = (p_A)(1 - c_A) + (1 - p_A)(-c_A)$
$EU_A(war) = p_A - p_Ac_A - c_A + p_Ac_A$
$EU_A(war) = p_A - c_A$

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 earns a payoff of $p_A - c_A$ if it fights. Instead, that value is state A's expectation for nature's move.

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

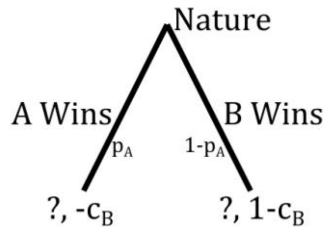

Here, state B loses and earns $-c_B$ with probability p_A, while it wins and earns $1 - c_B$ with probability $1 - p_A$. As an equation:

$EU_B(war) = (p_A)(-c_B) + (1 - p_A)(1 - c_B)$
$EU_B(war) = -p_Ac_B + 1 - c_B - p_A + p_Ac_B$
$EU_B(war) = 1 - p_A - c_B$

Thus, B earns $1 - p_A - c_B$ 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 outcome for B rejecting:

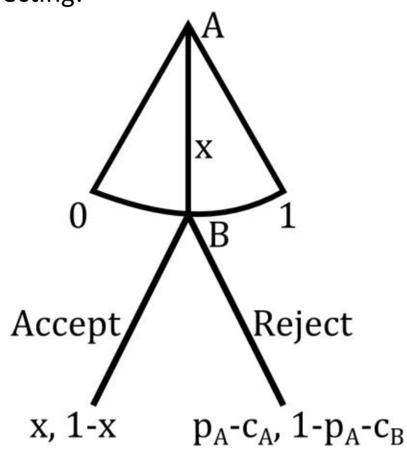

With this reduced game, we can now see which types of offers B is willing to accept. Let's focus on B's payoffs:

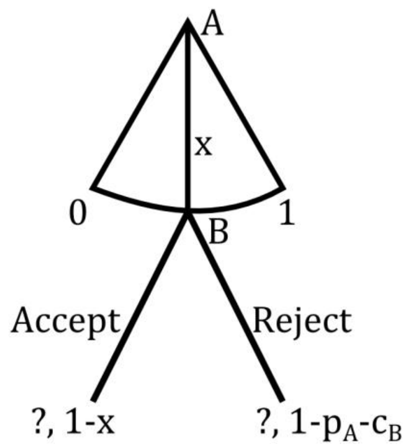

B can accept any offer 1 – x that is at least as good as $1 - p_A - c_B$, its expected utility for war. As an inequality:

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

Thus, B is willing to accept x as long as it is less than or equal to $p_A + c_B$.

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:

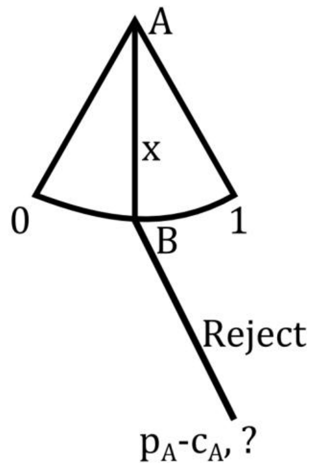

Consequently, we can bundle all of these scenarios into one expected utility. If state A makes an unacceptable offer to B—whether it is slightly unacceptable or extremely unacceptable—B always takes the same action, and A winds up with $p_A - c_A$.

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

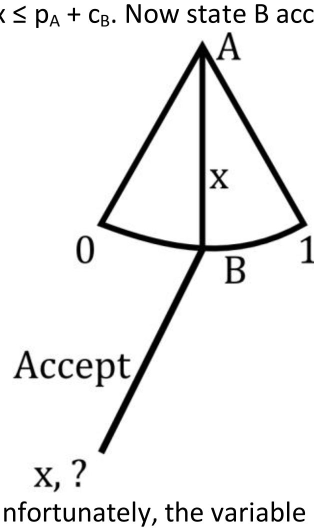

This time, state A simply earns x. Unfortunately, the variable payoff complicates matters. When B rejected, A earned the same payoff every time. Here, however, A's payoff is different for every acceptable offer it makes.

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 is willing to induce B to accept an offer, A wants that offer to be as beneficial as possible. Since B accepts any $x \leq p_A + c_B$, the largest x B is willing to accept is $x = p_A + c_B$. In turn, if A ultimately wants to make an acceptable offer, the best acceptable offer it can make is $x = p_A + c_B$; any smaller value for x needless gives more of the good to B.

Although we started with an infinite number of possible optimal offers (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 offer 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:

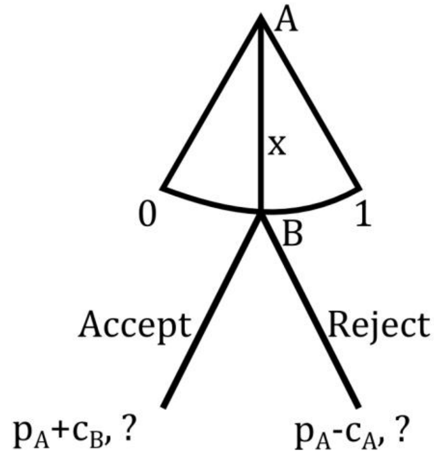

Thus, A can make the acceptable offer $x = p_A + c_B$ if its expected utility for inducing B to accept is greater than A's expected utility for inducing B to reject. As an inequality:

$p_A + c_B \geq p_A - c_A$
$c_B \geq -c_A$
$c_A + c_B \geq 0$

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

2.4: What Is the Puzzle?

All three of the above models indicate a range of practical alternatives to warfare always exists. However, as a matter of empirical fact, wars happen with disturbing regularity. *War's inefficiency puzzle* therefore asks why states chose to resolve their differences with inefficient fighting when they could simply select one of these alternatives instead.

A *rationalist explanation for war* responds to 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. We made some strong assumptions about the states knowledge of each other and the power structure over time. If we weaken these assumptions, the states may rationally end up fighting each other. In the following chapters, we will explore four of these explanations: preventive war, preemptive war, private information and incentives to misrepresent, and issue indivisibility.

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, power shifts over time. A weak country today can develop its economic base, produce more tanks at a later date, begin research into nuclear weapons, and become more threatening to its rivals at a later date. Thus, declining states might want to quash rising states before the latter becomes a problem in the future. Political scientists call this *preventive war* (or preventative war), and we cover it in the next chapter.

The baseline model also assumes that power remains static regardless of which state starts the war. If A starts the war, it wins with probability p_A; if B starts the war, A still wins with probability p_A. However, there may be certain advantages to striking first. After all, the initiator may benefit from surprising the other party and dictating where and when 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 fourth chapter.

Moving on, the states are perfectly aware of each other's military capabilities and resolve. 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 five explores such a possibility and shows how bluffing sabotages the bargaining process.

The sixth 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 Ken can, but they both cannot simultaneously be the king. Political scientists call this restriction *issue indivisibility* and express doubts whether it actually leads to war.