

National Security Strategy: Nash Equilibrium in Mixed Strategies

Professor Branislav L. Slantchev

January 13, 2014

Overview We have seen how to compute Nash equilibrium in pure strategies. We now learn how to compute equilibria in mixed strategies. We also learn how to calculate the probabilities of equilibrium outcomes and use them as the basis of a simple analysis of a generic crisis. We also look at an incident involving Greenpeace and conclude with an overview of classic strategic games.

1 The Crisis Game, Revisited

Recall that the crisis game, depicted in Figure 1, has two equilibria in pure strategies: $\langle E, \sim e \rangle$ and $\langle \sim E, e \rangle$. In these equilibria, the war outcome never occurs because one of the players submits. Of course, this begs the question why we thought the situation is a crisis in the first place: if one of the players was going to submit and both players knew that, then there really is no crisis. We now see that the game has another Nash equilibrium, this one in mixed strategies, that captures the idea of a crisis very well.




Figure 1: Crisis Game With Imperfect Information.

A pure strategy specifies what action to take at each information set where the player gets to move in the game. A mixed strategy specifies a probability distribution over the pure strategies. That is, it specifies the probability with which the player picks one of his pure strategies. In our crisis game, each player has one information set with two actions, so two pure strategies: $\{E, \sim E\}$ for player 1 and $\{e, \sim e\}$ for player 2. Let $p \in [0, 1]$ denote the probability with which player 1 chooses E , so $1 - p$ is the probability with which he chooses $\sim E$. Since p is a valid probability distribution, it is a mixed strategy. Since there is an infinite number of values that p can take, player 1 has an infinite number of mixed strategies. Analogously, let $q \in [0, 1]$ denote the probability with which player 2 chooses e , so that $1 - q$ is the probability with which she chooses $\sim e$. This q is a mixed strategy for player 2, and she has an infinite number of mixed strategies as well.

1.1 Best Responses and Mixed Strategies

We now must find the best responses given that players can use these mixed strategies. The principles are the same as for the pure strategies except we now must take into account the fact that outcomes are not certain but probabilistic. To see this, consider the following. Player 1 picks E with probability p . If player 2 responds by e , the game will end in disaster with that probability and with player 1's capitulation with probability $1 - p$. Since the outcome is uncertain, player 2 must compute the **expected utility** of her strategy. When the outcomes were certain, we simply compared the utilities attached to them to decide which action is better. We cannot

do this when the outcomes are uncertain because one action can result in more than one outcome. The Expected Utility Theory by von Neumann and Morgenstern (to which I referred earlier) tells us how to deal with such situations.

To compute the expected utility of an action, you take the payoff for an outcome it produces and multiply it by the probability with which this outcome will occur; you do this for each outcome that the action can produce, and then add the results. In our example, disaster occurs with probability p and yields a payoff of -5 for player 2 whereas capitulation by player 1 occurs with probability $1 - p$ and yields her payoff of 1. Player 2's expected utility of choosing e is:

$$EU_2(e) = p(-5) + (1 - p)(1) = 1 - 6p.$$

If player 1 chose E with probability $p = 1/4$, then player 2's expected utility from choosing e will be (substituting p in the expression above): $(1/4)(-5) + (3/4)(1) = -1/2$. In this way, we could compute the expected utility for any value of p .

Player 2's expected utility from playing $\sim e$ is computed analogously. Since player 1 chooses E with probability p , playing this strategy results either in player 2's capitulation (with probability p) or in the status quo (with probability $1 - p$). Her expected payoff is:

$$EU_2(\sim e) = p(-1) + (1 - p)(0) = -p.$$

The first term in the sum is the probability that player 1 chooses E multiplied by player 2's payoff from having to capitulate. The second term is the probability that he chooses $\sim E$ multiplied by her payoff from the resulting status quo. If player 1 uses $p = 1/4$, then player 2's expected utility from $\sim e$ would be: $(1/4)(-1) + (3/4)(0) = -1/4$.

Once you compute the expected utility for each strategy, the **best response is simply the strategy that yields the highest expected utility**. This is very similar to the best responses in pure strategies where we compared utilities directly. Under uncertainty, we have to compare expected utilities instead, that's all. For example, the best response to player 1's mixed strategy $p = 1/4$ is to choose the strategy $\sim e$ because it yields an expected payoff of $-1/4$, while choosing e yields an expected payoff of $-1/2$. This makes sense: given the very bad payoff associated with the disaster outcome, even a relatively low probability of it occurring should player 2 choose to escalate given player 1's strategy of escalating with probability $1/4$ keeps her from escalating.

This is the method for calculating the best responses to mixed strategies. Because there is an infinite number of mixed strategies, it is not possible to calculate all best responses directly. However, we can note something very important. Given any arbitrary mixed strategy p for player 1, player 2's best response would be e

whenever its expected utility exceeds the expected utility of $\sim e$:

$$\begin{aligned} \text{EU}_2(e) &> \text{EU}_2(\sim e) \\ p(-5) + (1 - p)(1) &> p(-1) + (1 - p)(0) \\ 1/5 &> p. \end{aligned}$$

Thus, we conclude that if player 1 chooses E with probability less than 20%, then player 2's best response must be to choose e . Conversely, we can flip the direction of the inequality to determine that if player 1 chooses E with probability greater than 20%, then player 2's best response must be to choose $\sim e$. Thus, even though player 1 has an infinite number of mixed strategies in which he plays E with probability less than $1/5$, player 2 has the same best response to all of them: she escalates. Analogously, even though he has an infinite number of mixed strategies in which he plays E with probability greater than $1/5$, she has the same best response to all of them: she does not escalate. Note that in both cases player 2's best response to player 1's mixed strategy is a pure strategy.

Since we have now covered all mixed strategies with $p < 1/5$ and all mixed strategies with $p > 1/5$, we have one remaining mixed strategy to consider: $p = 1/5$. When player 1 chooses this particular mixed strategy, player 2 is indifferent between escalating and not because her expected payoff from each are the same: $\text{EU}_2(e) = \text{EU}_2(\sim e) = -1/5$. Since both strategies are equally good (or equally bad), they are *both* best responses. This means that player 2 could use either strategy as a best response, but, more importantly, she can also choose to play either one of them with some probability; i.e., she can play a *mixed strategy* as a best response as well. This follows immediately from the fact that her payoff to the two pure strategies is the same: it does not matter which one she picks (and with what probability), she will always get the same payoff in expectation. Mathematically, if she uses e with probability q when $p = 1/5$, then her expected payoff from the mixed strategy q is

$$\text{EU}_2(q) = q\text{EU}_2(e) + (1 - q)\text{EU}_2(\sim e) = \text{EU}_2(\sim e) = \text{EU}_2(e)$$

regardless of the value of q . In other words, if player 1 mixes with probability $1/5$, then player 2 can do anything in response.

It is crucial to realize that **for a mixed strategy to be a best response, the player must be indifferent among the actions that this strategy uses**. If the player is not indifferent, then one of the actions must be yielding a higher expected utility than the other, but in this case a mixed strategy that assigns a positive probability to the action that yields a lower expected utility cannot be optimal. The player could do strictly better by choosing the action with the higher utility with certainty.

To see this, consider some $p > 1/5$. We already know that player 2's best response is the pure strategy $\sim e$ because $\text{EU}_2(\sim e) > \text{EU}_2(e)$ in this case. She is not indifferent between her two pure strategies, which means that any mixture between

them that puts positive probability on the strategy that yields the worse expected payoff is itself worse than the pure strategy with the better expected payoff. Intuitively, since the mixed strategy uses a bad pure strategy with positive probability, player 2 could do better by reducing the probability with which she chooses the bad pure strategy. In this case, this means playing $EU_2(e)$ with a lower probability, so she is always better off choosing q as low as she can. Since the lowest such probability is $q = 0$, she must be off simply not playing that pure strategy at all; i.e., she does best by playing $\sim e$ for sure, which is precisely the best response we found before. So remember, **if a mixed strategy is a best response, then all actions to which it assigns positive probability must yield the same expected utility to the player.** Furthermore, **if the player has more than one best response strategy to some strategy of the opponent, then any mixture among his best response strategies is also a best response to that strategy of the opponent.**

Let us now represent player 2's best responses in our mixed strategy notation. The pure strategy e can be represented by the mixed strategy in which she chooses e with certainty: $q = 1$. Analogously, the pure strategy $\sim e$ can be represented by the mixed strategy in which she chooses $\sim e$ with certainty: $1 - q = 1$, or $q = 0$. We can then summarize player 2's best responses in terms of q as follows:

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < 1/5 \\ q = 0 & \text{if } p > 1/5 \\ q \in [0, 1] & \text{if } p = 1/5. \end{cases}$$

Note that player 2's best response in each of the first two cases is a unique pure strategy. Only in the last case does she have an infinite number of best responses.

We can use a similar approach to determine player 1's best responses to player 2's mixed strategy q . Player 1 will choose E whenever:

$$\begin{aligned} EU_1(E) &> EU_1(\sim E) \\ q(-5) + (1 - q)(1) &> q(-1) + (1 - q)(0) \\ 1/5 &> q \end{aligned}$$

Conversely, he would choose $\sim E$ whenever $q > 1/5$, and will be indifferent between the two actions whenever $q = 1/5$. Summarizing his best responses in terms of p gives us:

$$BR_1(q) = \begin{cases} p = 1 & \text{if } q < 1/5 \\ p = 0 & \text{if } q > 1/5 \\ p \in [0, 1] & \text{if } q = 1/5. \end{cases}$$

These best responses are very intuitive: each player chooses to escalate if the probability that its opponent will escalate is sufficiently low (in this case, less than 20%); otherwise, the player prefers to submit.

1.2 The Mixed Strategy Equilibrium

Recall that an equilibrium is a strategy profile where all strategies are best responses to each other. Let's see which profiles in this game meet this criterion. Recall that a strategy profile in mixed strategies is denoted by $\langle p, q \rangle$, where p is player 1's strategy, and q is player 2's strategy.

First, consider strategy profiles with $p > 1/5$. In this case, player 2's best response is $q = 0$. The best response to $q = 0$ is $p = 1$. Therefore, the profile $\langle 1, 0 \rangle$ is a Nash equilibrium. This, of course, is the pure-strategy equilibrium $\langle E, \sim e \rangle$ we already know about.

Second, consider strategy profiles with $p < 1/5$. In this case, player 2's best response is $q = 1$. Player 1's best response to $q = 1$ is $p = 0$. Therefore, the profile $\langle 0, 1 \rangle$ is another Nash equilibrium. This, of course, is the pure-strategy equilibrium $\langle \sim E, e \rangle$ that we have already seen.

Finally, consider strategy profiles with $p = 1/5$. We know that if player 1 is willing to mix with this probability in equilibrium, he must be indifferent between his two pure strategies. From his best responses, we further know that he will only do so if player 2 chooses $q = 1/5$. Thus, in any strategy profile in which player 1 mixes with $p = 1/5$, it must be that player 2 mixes with $q = 1/5$. But since player 2 must be willing to mix in equilibrium, she must also be indifferent between her two pure strategy. From her best responses we know that she will only do so when $p = 1/5$. Thus, in any strategy profile in which player 2 mixes in equilibrium, it must be that $p = 1/5$. We conclude that the only mixed-strategy profile that can be an equilibrium involves both of them mixing with probability $1/5$. These mixed strategies are mutual best responses and the profile is the **mixed-strategy Nash equilibrium** (MSNE).

It is important to realize that when player 1 chooses $p = 1/5$, player 2 could use *any* mixed strategy as a best response. She has no particular reason to pick $q = 1/5$. However, if it is optimal for player 1 to mix (that is, if his strategy is part of an equilibrium), then he must be expecting player 2 to mix with $q = 1/5$ because otherwise he would choose the pure strategy that happens to be a best response to the strategy player 2 is playing. Therefore, in equilibrium player 2's strategy must be $q = 1/5$. It is crucial to note that she does not pick that strategy *in order* to make player 1 indifferent. Rather, it is because player 1 is indifferent in equilibrium that she must be playing (or at least he must be expecting her to play) this particular mixed strategy.

The MSNE of the crisis game now looks like a real crisis. In the solution of the game, escalation occurs with positive probability. Let's calculate the probabilities of the various outcomes. The probability of war is the probability that both escalate: $1/5 \times 1/5 = 1/25$, or 4%. The probability that 1 escalates and 2 submits is $1/5 \times 4/5 = 4/25$, or 16%. This is also the probability that 2 escalates and 1 submits. Finally, the probability that the status quo prevails is $4/5 \times 4/5 = 16/25$, or 64%. Let's check our

calculations: $64 + 16 + 16 + 4 = 100$. That is, the probability of one of these four outcomes occurring is 100%.¹

We can learn quite a bit from this solution. First, the probability that neither site escalates and the status quo prevails is rather large, 64%. This should be intuitive: since crises are dangerous games to play, most often than not players will avoid them. The status quo will have a strong pull and many would-be confrontations would simply never materialize because states would be afraid of the risks involved. We should keep this in mind when we study deterrence (the main security strategy of the US during the Cold War), and especially its claims that it has prevented a confrontation with the USSR. Contrast this with the finding from the PSNE: the SQ *never* survives in these equilibria, which I said was counter-intuitive. Now our intuition is confirmed by the MSNE.

Second, the probability that a potential crisis escalates into a real one but is resolved by the submission of one of the participants is 32%. Thus, a significant portion of potential crises will be resolved short of war. In fact, conditional on the crisis occurring, the probability that it will be resolved by the submission of one of the participants is huge: 89%. To calculate this probability, note that only $100 - 64 = 36\%$ of potential crises become actual ones (in the others neither side escalates). Of these 36, only 4 end in war. Thus, the probability of a resolution by submission is $\frac{32}{36}$, or 0.888, which rounds to 89%.

Third, the probability that one particular player will prevail is 16%. Unfortunately, it is precisely the possibility of this outcome that leads players to engage in risky behavior—they escalate with positive probability because they are hoping that the other side would submit. Of course, because war is so bad, the probability of escalation is not too high (hence the high likelihood that the SQ would obtain.) Still, the probability of war erupting from a potential crisis in equilibrium is strictly positive at 4%. Conditional on an escalation occurring, the probability of war becomes non-negligible at 11%. Thus, in a crisis a risk of war always exists, which is what makes the confrontation so dangerous in the first place. Thus, even this fairly simple model tells us quite a lot about crises. We shall see how more elaborate models tell us more.

This MSNE reveals what will turn out to be a fundamental problem in strategic interstate interaction. Recall that Figure 1 depicts a situation in which *war is the*

¹I should note that the equilibrium probabilities we've derived here depend on the exact numbers we have specified for the payoffs, so one should not take them as general probabilities or anything like that. They just illustrate interesting consequences from this setup. As a good thought experiment, try altering the payoffs to one of the actors such that war is extremely bad for him. See what probabilities you get. This actually can tell you quite a bit about which side will be more resolved in a crisis and which side you'd expect to back down more often on average. Note also that when we wrote the best responses in terms of the mixed strategies, we were able to find all equilibria, both in pure and mixed strategies. This suggests that it may be very useful to go directly to this step to do the analysis instead of finding first the pure-strategy equilibria and then trying to figure out the mixed-strategy ones.

worst outcome for both players and this fact is common knowledge. Because both players also like the status quo quite a bit (it's their second most-preferred outcome), we can call them "peace-loving." And yet... optimal rational behavior produces a situation in which there is a strictly positive probability that these actors will go to war with each other! This is a fundamental result that we shall see in much more complicated settings and it bears repeating: war is not the outcome of some evil men plotting each other's destruction. Rather, it is the unavoidable consequence of rational players trying to obtain the best possible outcome (capitulation of the opponent). Tragically, this may sometimes end up producing the worse possible outcome instead.

It is also worth noting that the outcome of the interaction is uncertain in the sense that any one of the four possible outcomes occurs with positive probability in the MSNE. This indeterminacy is a direct consequence of the behavior of the actors, and is not due to some environmental chance events. To make this distinction clearer, note that one's action may have uncertain consequences simply because of intervening random factors that are inherently unpredictable. For instance, if I send a boat with troops to conquer an island, a freak storm could capsize the boat drowning everyone on board. Suppose the troops are certain to conquer the island because they are much stronger than the limited opposition the islanders could put up. Still, the outcome of my action of sending the troops will be uncertain: victory if the troops make it to the island safely or defeat if they don't. I have no influence over the freak storm happening, so from my perspective, the expected utility of sending the troops is the payoff from victory times the probability of a safe landing plus the payoff of defeat times the probability of a storm. The point to note here is that the probability of a storm is a type of *environmental uncertainty*: it's a factor beyond the control of either player.

Now contrast this with the crisis game. Here, one's action also has uncertain consequences: escalation may lead to war if the opponent happens to escalate too or it may lead to victory if the opponent does not escalate. As we have seen, my expected utility from escalation is the payoff from war times the probability that the opponent escalates plus the payoff from victory times the probability that the opponent does not escalate. The probability that the opponent escalates is a type of *strategic uncertainty*: it is certainly within the ability of the opponent to control it. The point of this distinction is that strategic interaction sometimes can involve this type of uncertainty (induced by the randomization strategies) and it is very insidious because actors create it on purpose. Of course, it makes perfect sense that they would—when revealing your pure strategy to the opponent would lead to behavior that will hurt you, you would certainly try to confound the opponent's expectations. This, of course, is going to complicate one's own decision-making because now it has to take place under this uncertainty.

What's the upshot of all of this? Suppose players engage in the crisis game and the outcome is war. Looking back, this clearly was the worst thing for both

of them. But can we then conclude that players made mistakes? No—they were pursuing their optimal strategies that involve an irreducible risk of war. Hence, the possibility of war actually occurring is part of their best strategies and when war happens one cannot really say that it was because someone made a mistake. Of course, in retrospect each player would dearly wish to have chosen the other action. But hindsight is 20/20: choosing non-escalation with certainty before war occurs is simply not optimal. When we look at history and see a war which was clearly against the interests of both actors, can we conclude that the actors were stupid, evil, or irrational? No—we have now seen how intelligent peace-loving rational players may inadvertently create a situation in which they end up at war with each other. Interpreting history is a lot trickier than looking at events after the fact and then judging them in the light of the knowledge that they have occurred. That is, we *know* for sure that war happened but the participants in the crisis were *unsure* whether their actions would actually precipitate it. From their perspective, the risk was rational and the gamble was worth it (recall that the risk of war in MSNE is 4% and the probability of outright victory is 16%). We cannot substitute our knowledge that the bad outcome actually happened for the rational gamble actors made before it did. Hence, it is quite difficult to pass judgment on such decisions.

2 The Greenpeace Interview

We use models to discipline our thinking. As economist Paul Krugman said, models are often smarter than we are. They force us to think through issues that might be complicated, unpleasant, or both. Their conclusions, once understood, may compel us to part with deeply held beliefs. A person who understands the simplest model will reason in far more sophisticated ways than a person who knows thousands of facts and figures but who does not have the analytical framework to make sense of them. That's why we want the models.

To illustrate what I mean, here's a paraphrased (I am quoting from memory) interview between an NPR host and a high-ranking Greenpeace member that I heard on the radio. The idea was the Greenpeace was interested in preventing some really large ship from going to some place where it was going to do some really unpleasant things to the environment. So, Greenpeace's strategy was to send a bunch of activists in small boats that would get really close to the ship. The danger of sinking them would presumably force the captain to turn the ship around and avoid killing a bunch of innocent civilians. The following exchange then occurred between the NPR host and the Greenpeace guy:

NPR Host: Are you not concerned for the safety of your men? You are sending people into a situation where there's a really high risk of them dying.

Greenpeace: There is no risk. The ship will turn around and none of our men will be harmed.

Let's be charitable and *assume* that the captain does care about the safety and health of a bunch of activists who, one should note, have voluntarily put themselves in harm's way. Let's suppose that the captain does prefer to avoid sinking their boats. In fact, let's go ahead and be extremely generous: let such an event be the *worst* outcome for the captain. You already see how many assumptions we need just to get the argument stacked in Greenpeace's favor.

So, the captain would like to get the ship to its destination with Greenpeace backing down (shows company is tough and won't allow itself to be coerced) most to turning away without activists circling in boats to turning away under duress (shows that the company can be blackmailed) to collision and disaster.

The activists, on the other hand, strictly prefer that the ship turns away under duress (shows Greenpeace is effective and preserves environment) to ship turning away on its own (only saves environment) to their boats backing down and the ship getting through (shows Greenpeace is ineffective and is bad for environment) to actually dying in a collision.

So, imagine the situation: the ship is going full speed to its destination when the Greenpeace inflatable boat races toward it. They can either swerve or keep going. If neither swerves, they collide and disaster occurs. If the ship swerves, the activists gain in ecological protection and reputational enhancements. If the activists only swerve, the company gains in dumping its stuff and rubbing the activists' noses in it. If they both swerve, the status quo prevails. Who, if anyone, would blink?

As you can probably guess (or shouldn't guess but verify by setting up the preferences), this situation can be described with our crisis game, so our conclusions carry over very nicely. The main conclusion is that there is a significant real danger of the activists getting killed in this crisis. But this is precisely what the Greenpeace strategy is actually implicitly relying on: the threat of their activists dying. If there were no danger, the captain of the ship would just keep going, forcing them to swerve. So, for the threat to work, they have to accept a large enough risk that the captain would believe that he will kill them by continuing. Only then would he decide to swerve. But the problem, as we shall see later, is that by the time you have convinced yourself that they are not going to swerve, it may be too late to swerve yourself. The situation may thus end in disaster anyway.

At any rate, the statement by the Greenpeace person on NPR is complete nonsense because she refused to even dwell on the unpalatable aspects of the activists' own strategy. In fact, they were relying on the risk of dying to threaten the ship captain and compel him to change course. The unpleasant truth is that for this threat to work, there must be a real danger of dying. We shall have a lot more to say about things like this one later on.

3 The Prisoner's Dilemma, Revisited?

Can playing a mixed-strategy get you out of the mess with the DA? Recall that for each strategy of the other player, the best response was always the pure strategy to testify. As we know, for a mixed strategy to be a best response, it must be the case that all actions to which it assigns positive probability yield the same expected utility. However, in this case, whatever the other player does, the best response is always the same. Therefore, this action will yield the highest expected utility no matter what mixture that other player may use. Hence, a mixed strategy that assigns positive probability to not testifying cannot be a best response. We thus conclude that there can be no Nash equilibrium in mixed strategies. Unfortunately, the unique Nash equilibrium of this game is the one we found in pure strategies. In it, both players defect and rat on each other.

4 Putting It All Together: Generic Games

We have now seen several strategic games like Chicken, the Stag Hunt, and the Prisoner's Dilemma, and in all cases we used specific numbers to represent the payoffs. When the games involve no uncertainty either because of chance moves outside the players' control or because of mixed strategies, the precise values of the payoffs do not matter, only their ordinal ranking does. However, when the game does involve chance — as it must whenever some player uses a mixed strategy — then the cardinal values become important. Why it is so is a bit technical (if you take my game theory course, you will find out), but essentially it is because risky choices involve attitudes toward risks and the sizes of the payoffs loom large in those calculations. When I am running a 20% risk of disaster for an 80% chance of the other player capitulating, it certainly matters not merely that disaster is worse than him capitulating but also just how much worse it is. The worse it is, the less willing I become to take my chances. Von Neumann and Morgenstern's Expected Utility Theory in fact specifies the assumptions about preferences over risky choices we need to make in order to ensure that we can represent these preferences with numbers and calculate the resulting expected utilities.

This might seem technical, but it matters for us because we wish to use games without necessarily specifying the precise values of the payoffs. Instead, we would like to use variables to represent these payoffs, and then examine what happens as we change their values. Consider a generic two-player simultaneous-move game where each player has only two pure strategies. We can represent it in a 2-by-2 payoff matrix, as in Figure 2. The mnemonics for the variables are W for “war”, V for “victory”, D for “defeat”, and S for “status quo”.

We shall now see how varying the ordinal rankings among these variables yields all the games we have seen so far, and how we can glean some additional insights

		Player 2	
		e	$\sim e$
Player 1	E	W, W	V, D
	$\sim E$	D, V	S, S

Figure 2: The Generic Game.

from representing them in this form. First, however, we shall make a crucial assumption that we shall maintain more or less throughout all models that we are going to analyze: **we shall assume that our players are not war-loving and do not like defeat: they always prefer both the status quo and victory to either war or defeat.** In our notation, we are going to assume that

$$S > W, \quad S > D, \quad V > W, \quad V > D.$$

The only variation we shall allow is between the rankings of S and V — which we can think of as the strength of the incentive players have to take advantage of the cooperative behavior of the opponent (do they reward cooperation with restraint and obtain S or do they exploit it and obtain V ?), and the rankings of W and D — which we can think of as their fear of being exploited (do they prefer to let it happen and obtain D , or would they rather avoid it and obtain W ?).

We are making these assumptions because otherwise our insights will be superficial: it is not going to be very helpful if we found out that players go to war in equilibrium when they both value war the most. This is not to say that this cannot happen (sure it can!) but that the analysis is trivial (and one hardly needs all this complicated game-theoretic machinery to do it). It would be much more interesting if we found that players go to war in equilibrium even though war is their least-preferred outcome. If this happens, and we understand why it does (game-theory to the rescue!), then we will have a deeper understanding of the possible reasons for conflict. This understanding can then help us analyze actual crisis cases and go beyond surface assertions about the causes of some behavior or other. This is what we are going to be doing for the rest of the course, which is why we wish to make our models as useful and interesting as possible.

What can we say about this game? We know that $\langle E, e \rangle$ will be an equilibrium whenever $W > D$. Moreover, it will be the unique equilibrium if $V > S$ too. In other words, if the complete ordering is

$$V > S > W > D,$$

then the game is a **Prisoner's Dilemma** and its unique equilibrium yields the pay-offs that are second-worst for the players. When the fear of being exploited ($W > D$) combines with a desire to take advantage of the other player ($V > S$), then players will be unable to coordinate on a cooperative outcome regardless of the amount of communication they are allowed to engage in.

If, on the other hand, $S > V$, then $\langle \sim E, \sim e \rangle$ will be an equilibrium as well. When the ordering is

$$S > V > W > D,$$

then the game is a **Stag Hunt**, and it has two-pure strategy equilibria, with $\langle \sim E, \sim e \rangle$ being the one both players prefer (it, in fact, yields the highest possible payoff for each player), but where $\langle E, e \rangle$ is risk-dominant, making it more likely for the players to coordinate on that profile and obtain their next-to-worst payoffs. Thus, making the status quo more attractive — which eliminates the desire to take advantage of the other player — can help, but the resulting situation (which still has the fear of being exploited looming as the worst possible outcome) still presents players with a difficult dilemma where the outcome can be very dependent on the amount of trust they have for each other. In most circumstances, this trust will not be enough to overcome to fear, and players will again end up with their next-to-worst outcome.²

You might be tempted to conclude that perhaps it is the fear of being exploited that is causing the problem here, so let's suppose players do not have it ($D > W$) but that they still want to take advantage of each other $V > S$. The resulting preference ordering will be

$$V > S > D > W,$$

and you can verify that this makes this a **Game of Chicken**. The two pure-strategy Nash equilibria are $\langle E, \sim e \rangle$ and $\langle \sim E, e \rangle$ but we know that there is going to be another one in mixed strategies as well. To find it, let p and q be probabilities with which player 1 and player 2 escalate, respectively. The expected payoff for player 1 can be computed as follows:

$$\begin{aligned} EU_1(E) &= qW_1 + (1 - q)V_1 = V_1 - q(V_1 - W_1) \\ EU_1(\sim E) &= qD_1 + (1 - q)S_1 = S_1 - q(S_1 - D_1). \end{aligned}$$

(We are using subscripts on the payoffs to keep track of which player we're referring to.) We know that player 1 will only be willing to mix when indifferent between his pure strategies, so in the MSNE it must be the case that $EU_1(E) = EU_1(\sim E)$. Solving this tells us that player 1 will mix only when he thinks that player 2 is going to escalate with probability

$$q = \frac{V_1 - S_1}{V_1 - S_1 + D_1 - W_1}.$$

(Note that the preference ordering ensures that this is a valid probability; i.e., a number between 0 and 1.) We further conclude that whenever player 1 is mixing,

²In fact, the Stag Hunt, like the Chicken game, also has an equilibrium in mixed strategies. It is specified exactly in the same way as we shall do for the Chicken game, so there is no need to do it here.

player 2 must be mixing as well, which in turn pins down the precise probability with which she must expect player 1 to escalate, which we derive by setting $EU_2(e) = EU_2(\sim e)$, or:

$$p = \frac{V_2 - S_2}{V_2 - S_2 + D_2 - W_2}.$$

We already know that in the MSNE the probability of war is positive, but we can now say something more about the crisis. For example, we can now ask questions like: “What happens to the probability that player 1 escalates if player 2’s payoff from victory (V_2) increases?” Try answering this first without analyzing the model. You might reason as follows: well, since player 2’s payoff from victory is now larger than before and she can only get this outcome by escalating, she should be more willing to escalate. In other words, increasing the payoff for victory should make her more willing to take risks to achieve that outcome, so q should go up. But since this makes escalation more dangerous for player 1 and his payoffs have not changed, he should be less willing to escalate. Thus, the increase in the victory payoff for player 2 must mean that she is more likely to secure the prize without a fight, and that the overall likelihood of war is smaller.

The first surprise is that player 2 will not, in fact, escalate with a higher probability in equilibrium. As you can see from the expression above, q is entirely independent of V_2 . This is because in equilibrium her escalation probability reflects player 1’s expectations about her behavior that make *him* indifferent, and this calculation naturally only involves player 1’s payoffs. Since these have not changed, q will not change either.

But how can that be? Our intuition seems to demand that an increase in V_2 must have some effect on behavior... and it does, just not where you would first expect it. Consider player 1’s strategy. You can see that p is a function of V_2 , and you can easily verify that it is, in fact, strictly *increasing* in that value.³ In other words, increasing player 2’s payoff from victory must make player 1 *more* likely to escalate in equilibrium! What?!?! This just made matters even more confusing!

This, however, what being “in equilibrium” really means. It means that players must be willing to stick to their strategies. Initially, player 2 is indifferent and so willing to play the mixed strategy. When her payoff from victory increases and nothing else changes, however, she will no longer be willing to mix: the expected payoff from escalation given the probability that player 1 escalates will now be strictly greater than the expected payoff from not escalating, and as a result she would actually strictly prefer to escalate. But if she is going to escalate, then player

³Just take the derivative:

$$\frac{dp}{dV_2} = \frac{D_2 - W_2}{(V_2 - S_2 + D_2 - W_2)^2} > 0.$$

1 will no longer be willing to mix either. In other words, the strategies would no longer constitute an equilibrium. If player 2 is going to continue to mix in equilibrium, it must be that she continues to be indifferent after V_2 increases, and since none of the other payoffs have changed, the only way this can happen is if player 1's probability of escalation increases as well. Since this puts more weight on the war outcome, it decreases the expected payoff from escalation for player 2 even when V_2 goes up. Thus, if the mixed strategies are going to remain optimal, an increase in V_2 will be met with an increase in p .

In other words, our intuitive logic has some parts right (e.g., that increasing V_2 will make player 2 prefer escalation) but fails to consider the entire effect (e.g., what happens when you put this fact together with the requirement that players choose best responses). This is why simple intuition might sometimes prove quite misleading.

Finally, observe that since p goes up and q remains constant, an increase in V_2 also leads to an increase in the equilibrium probability of war, which is $\Pr(\text{War}) = pq$. Thus, an increase in the value for victory for one of the players makes the other one more aggressive, and it makes it more likely that they will end up fighting.

Analogous arguments establish that when a player's value for war increases, then the probability with which his opponent escalates in equilibrium must increase as well (p is increasing in W_2 just like q is increasing in W_1). This also seems counter-intuitive: a player's dislike of fighting decreases but as a result his opponent becomes more likely to escalate. The overall effect might be less surprising: the equilibrium probability of war increases.

Conversely, when a player's value for the status quo increases, then his opponent's probability of escalation must go down (p is decreasing in S_2). This is surprising when you recall that the opponent prefers to take advantage of such failures to escalate. The overall effect, however, might be what you expect: the equilibrium probability of war decreases. At least we obtain an unambiguous prediction: if one is interested in preserving peace, then making the status quo more valuable (or war more costly) is the way to go.

5 Coming Up Next...

We have two more things to do before we can analyze several important games that relate to our study of national security. First, we have to learn how to analyze dynamic games. As you may have noticed, Nash equilibrium takes the entire strategies as given and requires them to be best responses to each other. However, we shall see that this solution concept has a significant shortcoming because it ignores the fact that one player may move before the other, and that given that move the player's response may not be optimal. Again, Nash equilibrium only considers the entire strategies for optimality, not parts of them. We shall learn the solution

concept called *perfect equilibrium* that takes care of that.

Once we learn how to analyze dynamic games by finding their perfect equilibria, we only need to learn how to analyze games of incomplete information. We shall see that in that case, the solution must also incorporate the beliefs of the players because their strategies will be optimal conditional on the beliefs they have. This solution concept is called *sequential equilibrium*, and we shall learn how to compute it.

We shall then use all of this to analyze several very important games from which we shall learn the ideas of credible commitments, signaling, bargaining, and screening, which form the core of the theories of the use of force: deterrence and compellence. All of these we shall then apply in our study of history of the Cold War and after. Finally, we shall look at several current problems through the lens of history and analytical analysis to see whether we can form an opinion about national security strategies one may pursue in these circumstances.