National Security Strategy: Deterministic Optimal Behavior

Professor Branislav L. Slantchev

October 7, 2019

Overview We know what strategies are available to the players in a game in extensive form. How are players going to play this game optimally? That is, what is the best strategy a player could pursue in it? We learn how to select the best strategy from the set of available strategies. The idea is that a strategy must be a best response to the other player's strategy. If both players are using strategies that are best responses to each other, then these strategies form a Nash equilibrium, the basic solution concept in game theory.

We now know how to take any situation we wish to analyze and represent it with an extensive form game. What we want to know is how rational players would then "solve" the game. That is, what actions are they going to take in pursuit of their interests? The idea here is that once we analyze the game and figure out what the optimal behavior (the solution) is, then we can extrapolate back to the real world and think about the original situation in those terms. In other words, we are trying to reduce the complexity of the real world to manageable proportions, analyze that stylized representation of it, and then see what we can learn about its underlying dynamics. This will produce insights about how players might be expected to behave, and what parameters are going to affect their behavior. We shall then take back these insights to the real world, and apply them to read history in a more analytically rigorous (and therefore more fruitful) way. We shall then combine our insights with what he have learned from history to current problems of national security.

So, the question before us is this: What actions should we expect players to take in any given game? On the most basic level, we would expect that players pursue their interests to the best of their ability. That is, they take an action that maximizes the likelihood of getting the best available outcome. The problem, of course, is that the outcome is determined jointly with the other player, who might be trying to obtain another outcome that she likes. To put it in our Cold War terms, the US may be trying to get the USSR to capitulate while the USSR may be trying to do the same to the US. So what strategies would the players use?

Intuitively, players will try to do their best. That is, they pick a strategy that will be the best course of action given what their opponents are doing. This is why we needed to list strategies: we need to (a) pick the best one among them, and (b) determine which the best ones are by looking at whether they are optimal responses to the other player's strategy.

1 Best Responses

The most fundamental concept in thinking about optimal (best) strategies is the idea of a **best response**. A "best response strategy" is a strategy that produces the most preferred outcome for the player *given* some strategy of its opponent. That is, you would take the strategy of the other player as fixed and then find your best response to it. The idea is that players would best-respond when they play the game. That is, they will choose a strategy that is a best response to their opponent's strategy. Any strategy that is not a best response cannot be optimal, so we should not expect players to choose it.

Let's now illustrate these ideas with our favorite simple crisis game with imperfect information, as shown in Figure 1. This time, instead of the labels US and USSR, we shall just call the players by their numbers.

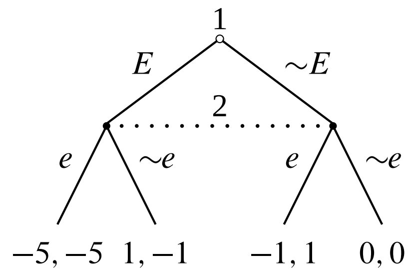

Figure 1: Crisis Game With Imperfect Information.

Player 1 has two pure strategies, $S_1 = \{E, \sim E\}$, and so does player 2: $S_2 = \{e, \sim e\}$. We shall use s_1 to denote an arbitrary pure strategy for player 1. Similarly, we shall use s_2 to denote an arbitrary pure strategy for player 2. Let's find the best responses for both players.

We begin with player 1. What are the best response strategies that it has? To find the set of best response strategies, we have to analyze all strategies of player 2. We begin with $s_2 = e$: What is the best response to this strategy? If player 1 chooses the strategy E, then the outcome will be disaster, and its payoff will be -5. If, on the other hand, player 1 chooses $\sim E$, the outcome will be victory for player 2, and the payoff for player 1 will be -1. Because $-1 > -5$, player 1 prefers to choose $\sim E$. That is, the strategy $\sim E$ is a best response to player 2's strategy e.

What is the best response to $\sim e$? Playing E produces victory by player 1 with a payoff of 1. Playing $\sim E$ produces the status quo outcome, with a payoff of 0. Because $1 > 0$, the best response to this strategy is to play E.

Hence, the best responses by player 1 are:

$$\mathrm{BR}_1(s_2) = \begin{cases} E & \text{if} \quad s_2 = \sim e, \text{ and} \\ \sim E & \text{if} \quad s_2 = e. \end{cases}$$

Best responses are always defined in terms of the opponent's strategy. That is, a best response is usually meaningless without reference to the other player's strategy. This is because a strategy must be a best response to *some* strategy of the opponent.

Let's now look at the best responses for player 2. Again, we have to consider all possible strategies for player 1. What is the best response to E? Playing e produces disaster, with a payoff of -5. Playing $\sim e$ produces capitulation by player 2, with the better payoff of -1, so $\sim e$ is the best response to E.

What about the best response to $\sim E$? Playing e produces victory for player 2 with a payoff of 1. Choosing $\sim e$ produces the status quo, with a payoff of 0, which is strictly worse. Therefore, the best response to $\sim E$ is e.

Hence, the best responses by player 2 are:

$$\mathrm{BR}_2(s_1) = \begin{cases} e & \text{if} \quad s_1 = \sim E, \text{ and} \\ \sim e & \text{if} \quad s_1 = E. \end{cases}$$

Notice (again) that the best response of a player is function of a strategy of its opponent.

What should we expect players to do? That is, how do we solve the game? A solution to the game must specify how both players are going to play it. In our example, a solution will consist of two strategies, one for player 1 and another for player 2. Such sets of strategies, with one strategy for each player, are called **strategy profiles**. For example, one possible strategy profile in our game is $\langle E, e\rangle$. In this strategy profile, player 1's strategy is E, and player 2's strategy is e. Obviously, the number of such (pure-strategy) profiles can be obtained by multiplying the number of pure strategies for player 1 by the number of pure strategies for player 2: $2 \times 4 = 8$ strategy profiles. Which of these strategy profiles are solutions to the game?

2 Nash Equilibrium in Pure Strategies

Let's think about what it is that we are trying to achieve here. We want to find the optimal ways to play this game. Clearly then, any solution must involve best responses. That is, in any solution each player must be playing one of his best response strategies. But, as we have seen, whether a strategy is a best response depends on the strategy of the opponent. For example, the strategy $\sim E$ is a best response if player 2 is playing e but it is not a best response if player 2 is playing $\sim e$. Thus, the strategy profile $\langle \sim E, \sim e\rangle$ cannot be a solution because player 1's strategy is not a best response to player 2's strategy (and neither is player 2's strategy a best response to player 1's).

But if some player's strategy that is included in the profile is not a best response, then we cannot consider this profile to be a solution because that player would not play this strategy but will instead pick one that is a best response. Thus, any solution to the game must specify best-responses for all players, but (b) these responses depend on the strategy of the opponent. The conclusion suggests itself: A solution must include only strategies that are best responses to each other. This ensures that all strategies in the profile are best responses. We shall consider every strategy profile of this kind a solution to the game, and will call it an **equilibrium**.

We now have our fundamental *solution concept* for the analysis of games, the idea of an equilibrium. It is named in honor of John Nash, who invented it in the 1950s. **A Nash equilibrium is strategy profile with the property that the strategies for all players are best responses to each other.** When we analyze the games, we shall look for their equilibria. Because in equilibrium all players are using best-response strategies, **in a Nash equilibrium, no player can profit by changing its strategy given the strategies of the other players.** This is why these strategy profiles are called "equilibria": everyone is behaving optimally, and no one wants to change anything unilaterally.

Let's find the equilibria in our example. We already know that $\langle \sim E, \sim e \rangle$ is not an equilibrium because it the strategies are not best responses to each other. Let's examine the remaining strategy profiles by looking at the definitions $\text{BR}_1(s_2)$ and $\text{BR}_2(s_1)$. Consider first strategy E. From $\text{BR}_2(E)$, we know that it has one best responses: $\sim e$. Is E itself a best response to $\sim e$? Sure it is: $\text{BR}_1(\sim e) = E$. So, one Nash equilibrium is $\langle E, \sim e \rangle$.

How do we interpret it? The Nash equilibrium $\langle E, \sim e \rangle$ is telling us that if player 1 thinks that player 2 is going to play $\sim e$, then its optimal response is to play E. Conversely, if player 2 thinks that player 1 is going to play E, then its optimal response is $\sim e$. The strategies are mutually optimal. No player has an incentive to deviate (switch to another strategy) given its opponent's strategy.

What does the solution tell us? It allows us to predict an outcome of the game. Because a profile specifies one strategy for each player and because a strategy is a complete plan of action, each profile produces an outcome for the game. This is the outcome that will result when players implement the actions specified by the strategies in this profile. In our Nash equilibrium profile, the outcome follows escalation by player 1 and submission by player 2, that is, capitulation by player 2.

So, a solution tells us both what strategies are optimal and what we should expect the outcome of the game to be. An outcome produced by an equilibrium profile is called **equilibrium outcome**.

Let's continue with our examination of the strategy profiles. Are there any other solutions? Consider now $\sim E$. We know that $\text{BR}_2(\sim E) = e$. Is $\sim E$ itself a best response to e? Sure it is: $\text{BR}_1(e) = \sim E$. Hence, the profile $\langle \sim E, e \rangle$ is another Nash equilibrium. The see the outcome produced by this strategy profile, note that in it player 1 submits, and player 2 escalates. The equilibrium outcome is victory by player 2.

In this equilibrium, player 1 expects that player 2 is going to escalate. Given such a strategy, player 1's optimal course of action is to submit because even though this would result in capitulation, it would avoid the worst outcome of war. Given that player 1 is expected to submit, player 2's strategy of escalating is also optimal.

Remember that *to establish that a profile is not an equilibrium, all we have to do is find* ***one*** *player who could improve his payoff by switching to some strategy other than the one prescribed by the strategy profile.* That is, establish that a player has a profitable deviation while holding the strategy the other player constant as prescribed by the profile. So, to eliminate $\langle E, e \rangle$, we could either show that player 1's best response to e is $\sim E$, **or** show that player 2 can improve her payoff by switching to $\sim e$ when she expects player 1 to choose E. We don't have to do both.

On the other hand, *to show that a profile* ***is*** *an equilibrium, we have to verify that* ***each*** *player's strategy is a best response to the strategy of the other player.* Thus, showing that something is an equilibrium is more involved than eliminating a strategy profile as not being one.

We conclude that our crisis game of imperfect information has two Nash equilibria in pure strategies, $\langle E, \sim e \rangle$, and $\langle \sim E, e \rangle$. That is, the game has two solutions in pure strategies. We shall see next time that it also has a solution in mixed strategies that can tell us more about the dynamics of the game than these two pure-strategy equilibria. Right now, we've learned that although two solutions are possible, we don't really know which one the players are going to end up playing. In fact, this is one of the defining features of a crisis: opponents simply do not know what the outcome will be, mostly because they are not sure about what the other player is going to do.

Each solution tells us that *if players expect to play the equilibrium strategies*, then nobody would have an incentive to deviate. So, the actual equilibrium chosen depends on the beliefs of the players. We shall have more to say about this later.

We can, however, already say a few interesting things based on our analysis. First, none of the equilibrium profiles involves war. That is, we can say that rational players will never go to war with certainty if they find themselves in the situation described by the game. This should not be surprising—after all, war is the worst possible outcome for both of them. Second, none of the equilibrium profiles preserves the status quo. That is, we can say that rational players will not be able to keep the status quo with certainty if they find themselves in the situation described by the game. This is now much more disturbing—recall that the status quo is the second most preferred outcome for both players—and yet rational play implies that it is not certain that the status quo would survive the interaction. In either of the two profiles, the status quo will be altered to the advantage of one of the players. (As we shall see next time, this game will also produce a positive risk of war in the one other equilibrium that we will find when we consider mixed strategies.) Hence, from our perspective, we can see that whereas war is not certain, neither is peace without some changes. In that sense, peace is unstable even though both actors dislike war and like the status quo only second to victory over the opponent.

3 Two Social Dilemmas and Two Takes on Arms Races

3.1 The Prisoner's Dilemma

You and that no-good friend of yours your mother always warned you about get busted by the cops. You are immediately placed in separate cells. The DA comes to you and gives you the following spiel:

> *We have videotape evidence showing that you two shoplifted CDs from Borders. We can easily convict you right now to 10 months in jail each. However, we strongly suspect that you two are behind a string of CD disappearances in the area. If you agree to testify against your partner, we'll let you go unless he testifies as well, in which case we shall have enough evidence for a serious sentence, but because you cooperated*

> *we'll reduce it to 15 months. If you refuse to testify against him, then you will get your 10 months sentence if he also keeps silent, but you will get the full 3 years sentence if he testifies against you. I have told the same thing to the other guy. Here's a piece of paper, where you can write your testimony against him.*

Do you testify against the other guy?

Well it depends on what you think he is going to do. Clearly, the best outcome for both of you is to keep quiet and get the 10 months sentence. But you can't communicate to arrange that. So, what is the optimal strategy here?

We set up the situation as a two-player game: you and your partner are the players. Each of you has two strategies: cooperate with the other prisoner and do not testify, C, or defect and testify, D. There are four possible outcomes: both cooperate so neither testifies (CD Theft Conviction), you defect and testify but he cooperates and keeps quiet (Freedom for you, Shafted for him), he defects and testifies but you cooperate and keep quiet (Shafted for you, Freedom for him), and both defect and testify (Multiple Thefts Conviction). The payoffs (in terms of months you get to spend in jail) are given in Table 1, which also lists the preference ordering over the various outcomes specified in terms of strategy profiles where you are the first player.

	Freedom	$\succ$	CD Theft	$\succ$	Multiple Thefts	$\succ$	Shafted
	0	$>$	-10	$>$	-15	$>$	-36
You	(D,C)	$\succ$	(C,C)	$\succ$	(D,D)	$\succ$	(C,D)
Your Partner	(C,D)	$\succ$	(C,C)	$\succ$	(D,D)	$\succ$	(D,C)

Table 1: Payoffs in the Prisoner's Dilemma.

Since each is deciding what to do in ignorance of what the other one is doing, this is a game of imperfect information. Figure 2 gives the full representation of this game that we are going to analyze.

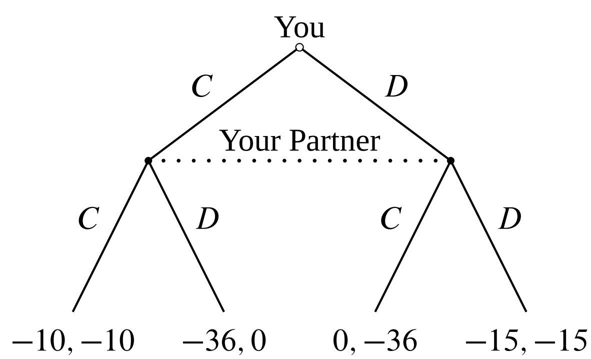

Figure 2: The Prisoner's Dilemma.

Let's determine your best responses. If the other guy is going to testify, you would get -15 if you testify as well, and -36 if you keep quiet. Your best response is to testify. If the other guy is going to stay quiet, you would get 0 if you testify, and -10 if you keep quiet. Your best response again is to testify:

$$\text{BR}_{you}(D) = \text{BR}_{you}(C) = D.$$

The situation is symmetric, so following the same line of logic, you get your partner's best responses:

$$\text{BR}_{he}(D) = \text{BR}_{he}(C) = D.$$

This means that there really is only one strategy profile that is a Nash equilibrium candidate because the only possible combination of best responses is the profile $\langle D, D \rangle$. Is it an equilibrium? Sure it is: the best response to D is D for both players.

Thus, we conclude that this game has a unique Nash equilibrium, in which both players defect and confess. This tells you that if you and your partner are both rational, each should expect the other to rat on him, and should expect the other to expect him to rat on his partner, and so on. You better start scribbling on that piece of paper. The DA, of course, is well aware of this, which is why you were placed in this situation. The equilibrium outcome is that both prisoners sing like nightingales and each gets 15 months in their local friendly county jail.

The Prisoner's Dilemma illustrates (in a very simple setting) a profound result that we shall see numerous times throughout this course as it applies to international relations. Observe that *both* players would be better off if they cooperate and keep quiet. They also *know* that this is the case. And yet, when they rationally pursue their best interests, they end up with an outcome that is *worse* for both of them.

That, in essence, is the tragedy of interstate politics: actors may be able to identify an outcome that they all would rather have, they may all agree on what this outcome is and what they need to do in order to get there, and yet they may still fail to obtain it because the individually rational actions they take produces an outcome that is socially suboptimal for all. Hence, as it often happens in world politics, the problem is not that people are too dumb to see the solutions or even that they strenuously disagree about them. The problem is that these solutions often require actions that are not in the best interest of some player given what the others are doing. If that's the case, players will behave according to their preferences and, unfortunately, the outcome will be worse for everyone involved. Very often, then, we shall see that interstate politics is the search for second-best (if that) solutions, with the very best sadly remaining forever out of reach because of these "perverse" strategic incentives of the actors involved.

3.2 The Arms Race as a Prisoner's Dilemma

The Prisoner's Dilemma is an abstract description of many social situations. In fact, to go back to national security, you can think of it as a description of a simple arms race. Suppose it is determined that a new technology has just emerged and that it allows both us and our enemy to produce a super weapon that can guarantee winning a confrontation against an opponent who does not have it. The confrontation is very important. If both have the weapon, the effects cancel each other out. It takes a year to construct the weapon, but once built, it becomes immediately useful. The weapon is quite costly and each nation must shift resources from consumer goods to the military sector, which is politically unattractive. Should we build the weapon or not?

We have already simplified the situation drastically in this description. Let's now represent it with a game. There are two players, "us" and "they." Each has two options: defect and build the weapon, D, or cooperate and do not build it, C. There are four outcomes: both build the weapon (an arms race), only one builds the weapon (the one that does wins), or neither does (status quo). What are the preferences over these outcomes? If we arm, we pay the cost of doing so regardless of whether the weapon is used or not. Assume that the confrontation is very important, and the benefits of winning it exceed the costs of producing the weapon.

If only the enemy arms, we don't pay the cost of arming but lose the confrontation, which is really bad: *defeat*. If we arm and the enemy arms as well, then we pay the cost but since nobody can get the upper hand, no confrontation occurs: *arms race*. If we are the one side with the weapon, then we pay the cost but win the confrontation, which is really good: *victory*. If neither side arms, no confrontation occurs: *status quo*. Thus, we have the following preferences:

$$\text{Victory} \succ \text{Status Quo} \succ \text{Arms Race} \succ \text{Defeat}.$$

Note that victory is preferred to the status quo because the benefits from winning the confrontation are so high that even when we factor in the costs of building the weapon, it is still better than the status quo life with the enemy. The status quo, however, is preferred to an arms race because with an arms race we pay the costs of building the weapon but we don't get anything out of it except that the enemy can't defeat us, which is what the status quo already is. Finally, the arms race is preferred to defeat because losing is so disastrous that it is worse than avoiding the costs of building the weapon. Since the situation is symmetrical, our opponent has similar preferences.

We can represent these preferences in many ways, as we already know. Table 2 lists one possibility. You can think of the first row as follows: the cost of the weapon is \$5 million. The prize of a confrontation is \$15 million. The status quo (no weapon, no prize) is then worth 0. Victory is worth the prize minus the costs: $15 - 5 = 10$ million. The arms race (weapon, no prize) is worth the costs of the weapon: -5 million. Finally, defeat (no weapon, lost prize) is worth losing the prize: -15 million.

The table also lists the preference ordering over the strategy profiles (outcomes) with us as the first player. Clearly, these preferences are identical to the Prisoner's

	Victory	$\succ$	Status Quo	$\succ$	Arms Race	$\succ$	Defeat
	10	$>$	0	$>$	-5	$>$	-15
Us	(D, C)	$\succ$	(C, C)	$\succ$	(D, D)	$\succ$	(C, D)
Them	(C, D)	$\succ$	(C, C)	$\succ$	(D, D)	$\succ$	(D, C)

Table 2: Payoffs in the Arms Race Game as a PD.

Dilemma. This means that the Arms Race game is strategically equivalent to the Prisoner's Dilemma. We don't even have to do a separate analysis here. From the solution of the former, we know that this game has a unique Nash equilibrium, and it is $\langle D, D\rangle$. The equilibrium outcome is an arms race: both players lose because they pay the costs of building the weapons but do not get any benefit from having them.

3.3 The Stag Hunt

The Prisoner's Dilemma is one type of social problem which assumes that unilateral defection is preferable to mutual cooperation. There are, however, situations in which mutual cooperation is the most preferred outcome for both players. And yet, as we shall now see, this in no way guarantees their ability to cooperate!

The classic illustration of such a social dilemma is due to Jean-Jacques Rousseau, and the story goes as follows. Two hunters must decide whether to cooperate, C, and hunt a stag together, or defect, D, and chase after a rabbit individually. If the both stalk the stag, they are certain to catch it, and they can feast on it. However, it requires both of them to stalk it, and if even one of them does not, the stag is certain to get away. If, on the other hand, a hunter goes chasing a rabbit, he is certain to catch one regardless of what the other one does. Assume that if the other one is also hunting for rabbits, the noise they both make scares the tastiest rabbits away and they can only catch stale hares with lower nutritional value. In other words, if you go after a rabbit, there is a slight preference that you do so on your own. Even the best rabbit is worse for a hunter than his share of the stag. There is only time to stalk the stag or hunt for rabbits, they cannot do both. You are one of these hunters. What do you do?

We set up the situation as a two-player game: you and the other hunter are the players. Each of you has two strategies: cooperate, C, or defect, D. There are four possible outcomes: both cooperate and catch the stag (Stag), you chase a rabbit and he stalks the stag (Tasty Rabbit for you, Hunger for him), you both hunt for rabbits (Stale Hare), and you stalk the stag while he catches a rabbit (Hunger for you, Tasty Rabbit for him). One possible specification of the payoffs that reflects the preferences is given in Table 3, which also rank orders the outcomes represented by the strategy profiles in which you are the first player.

Compare the rankings of the strategy profiles to those in the Prisoner's Dilemma

	Stag	$\succ$	Tasty Rabbit	$\succ$	Stale Hare	$\succ$	Hunger
	4	$>$	3	$>$	2	$>$	0
You	(C, C)	$\succ$	(D, C)	$\succ$	(D, D)	$\succ$	(C, D)
Other Hunter	(C, C)	$\succ$	(C, D)	$\succ$	(D, D)	$\succ$	(D, C)

Table 3: Payoffs in the Stag Hunt.

shown in Table 1. As before, unreciprocated cooperation is the worst possible outcome for each player, and mutual defection is the second worst outcome. Unlike the PD, however, the preferences in a Stag Hunt situation are such that both players prefer mutual cooperation to unilateral defection.

Since each is deciding what to do in ignorance of what the other one is doing, this is a game of imperfect information. Figure 3 gives the full representation of this game that we are going to analyze.

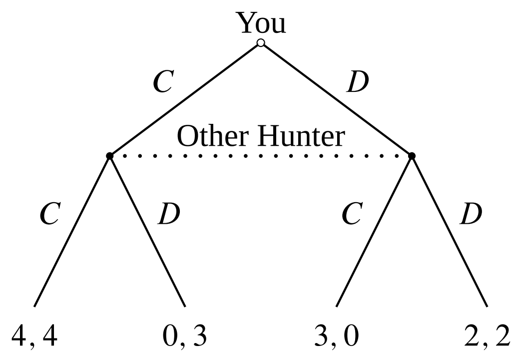

Figure 3: The Stag Hunt.

Let's determine your best responses. If the other hunter is going to stalk the stag, you would get the stag if you cooperate as well, and you would get the juicy rabbit if you defect. Since the stag payoff of 3 is better than the juicy stag payoff of 2, your best response is to cooperate. If the other hunter is going after a rabbit, then trying to cooperate would just leave you hungry (with a payoff of 0), whereas chasing a rabbit would at least guarantee you a stale hare (payoff of 1). Your best response is to defect and chase the hare. Putting this all together yields:

$$\mathrm{BR}_{you}(C) = C \quad \text{and} \quad \mathrm{BR}_{you}(D) = D.$$

The situation is symmetric, so following the same line of logic, you get the other hunter's best responses:

$$\mathrm{BR}_{he}(C) = C \quad \text{and} \quad \mathrm{BR}_{he}(D) = D.$$

Examination of the four strategy profiles shows that there are two in which the strategies are mutual best responses. Both $\langle C, C\rangle$ and $\langle D, D\rangle$ are Nash equilibria in pure strategies.

Unlike the Prisoner's Dilemma, mutual cooperation can be sustained in equilibrium. Unfortunately, like the Prisoner's Dilemma, mutual defection can also be an equilibrium. In that sense, assuming that both players prefer mutual cooperation to every other possible outcome does not actually mean that they will cooperate! This is a fairly startling result and it is worth thinking through why it happens.

Recall that a best response is a strategy that is optimal given what you think the other player is doing. In this sense, cooperation is best if you think the other is cooperating. In a Nash equilibrium, these expectations are self-enforcing in the sense that *your* expectation of the other player choosing to cooperate rationalizes *your* choice to cooperate, which in turn validates *his* expectation that you will cooperate, which then rationalizes *his* choice to cooperate, and this in turn validates *your* expectation that he will cooperate, closing the circle of mutually supporting expectations.

Unfortunately, the exact same logic applies in the case of defection. If you think your partner will defect, you will defect as well, which validates his expectation that you will defect, which rationalizes his defection, which in turn validates your expectation that he will defect. Again, the circle is complete and we have an equilibrium with mutually supporting expectations.

The question then seems to boil down to where we "begin" the circle of expectations. For instance, if we think one of the hunters expects the other to cooperate, we end up with the cooperative equilibrium. If, on the other hand, we think of the hunters expects the other to defect, we end up with the non-cooperative equilibrium. So which expectation is more likely? Without knowing the hunters and their relationship, it is impossible to say for sure. However, we could ask ourselves: if I were one of these hunters, which is the *least risky* choice to make? That is, which choice gives me an outcome that leaves me least vulnerable to the behavior of the other hunter?

In a sense, we are trying to protect ourselves from a mistaken expectation. Let's say I generally trust the other hunter to cooperate but I also know that sometimes he gets tempted when he sees rabbits, and I am not entirely sure that he will not see a rabbit or that if he sees one while stalking the stag, he won't abandon the stalking in order to chase after the rabbit. Now, if I cooperate, I would get the stag if he does not get distracted but I will end up hungry if he does. If I defect, I would get the juicy rabbit if does not get distracted, and I will end up with a stale hare if he does. When I cooperate, the worst possible thing that can happen to me is to go hungry. When I defect, the worst possible thing that can happen to me is to end up with a stale hare. In that sense, defection is less risky because it leaves me less vulnerable in the case that I have misjudged my partner or he makes a mistake.

In case you are wondering, this can be formalized precisely. The notion of *risk-dominance* is due to Harsanyi and Selten, and for this game it can applied as follows. For each equilibrium, we can compute the product of losses if someone deviates from it. Consider your situation first. You are supposed to play the cooperative equilibrium $\langle C, C\rangle$ but instead you deviate it. Since C is a best response to C, this deviation is going to cost you: your payoff from $\langle D, C\rangle$ cannot exceed the equilibrium payoff by the very definition of equilibrium. In this case, you are going to suffer a *deviation loss* of $4-3 = 1$. Consider now the non-cooperative equilibrium $\langle D, D\rangle$ and suppose you deviate from your strategy. This time, you will end up at $\langle C, D\rangle$ with a deviation loss of $2 - 0 = 2$. Compare now your two deviation losses: since the loss from $\langle D, D\rangle$ is greater than the loss from $\langle C, C\rangle$, you should be *less likely* to deviate from $\langle D, D\rangle$. Intuitively, you stand to lose more if you do so, so you would have less incentive to do it. From the other player's perspective, then, $\langle D, D\rangle$ appears less risky: you are more likely to stick with the equilibrium strategy. We can now apply the same argument to the other player, and since her deviation loss from $\langle D, D\rangle$ exceeds her deviation loss from $\langle C, C\rangle$, it makes sense that you should consider it more likely that she should stick with her equilibrium strategy under $\langle D, D\rangle$.

Putting these two together, we can compute the *risk-dominance* of one equilibrium profile over another. Take the product of the deviation losses for the players: for $\langle C, C\rangle$ it is $1 \times 1 = 1$, whereas for $\langle D, D\rangle$ it is $2 \times 2 = 4$. The profile with the *higher* product of losses is said to be risk-dominant: it is the one that players are more likely to stick with. In this game, the risk-dominant profile is $\langle D, D\rangle$, and as a result players should be more likely to expect $\langle D, D\rangle$ than $\langle C, C\rangle$. As a result, we would expect $\langle D, D\rangle$ to be the equilibrium they coordinate on, mutual defection will be the outcome.

You can also arrive at the same conclusion with slightly different reasoning. When would you choose C over D if you are not entirely certain what the other hunter would do? If you choose C, you would get 4 if she also chooses C, and you would get 0 if she chooses D. Let p be the probability with which she chooses C, so $1 - p$ is the probability with which she chooses D. Then, your expected payoff from C is just $4p + 0(1 - p) = 4p$. If you choose D, you would get 3 if the other hunter chooses C, and you would get 2 if she chooses D as well. Your expected payoff from D is then $3p + 2(1 - p) = p + 2$. You would choose C when $4p > p + 2$, or when $p > 2/3$. In other words, you would have to believe that the probability that the other hunter is going to cooperate is *at least* ⅔, or (66%), before you would want to cooperate too. An analogous computation for her reveals that she has to believe that you must cooperate with probability at least ⅔ before cooperation would be rational for her. The *risk-factor* for $\langle C, C\rangle$ is ⅔.[1] In contrast, defection is rational if each of us believes that the chance that the other will defect is at least ⅓. Since the risk-factor of $\langle C, C\rangle$ is higher, we would expect players to choose $\langle D, D\rangle$.

The risk-dominance argument would select the non-cooperative equilibrium even

[1]In this particular game, the minimum probability that one player has to assign to the other playing the equilibrium strategy is the same for both. If they are different, the risk-factor for the equilibrium is the smallest of the two.

though one might initially believe that rational actors would surely coordinate on the cooperative one: after all, both of them would get better payoffs in $\langle C, C\rangle$ than they do in $\langle D, D\rangle$. In the context of a stag hunt, the advantage of avoiding the worst-case scenario might not be obvious, at least not as obvious as it is when we recast the Stag Hunt as an arms race (which we shall shortly do).[2]

Even small doubts about his trustworthiness may make me think about defection. Now, it gets worse if you consider what this means for my partner. Suppose he is aware that I harbor small doubts about his ability to resist temptation. Suppose he is resolved to resist it too. The problem is that when he is aware of my doubt, he knows that I may be tempted to protect myself to avoid going home hungry. But this then makes him even more tempted to defect in order to protect himself from being left with nothing. And of course, I am aware of all of this, which makes me even more suspicious that he might actually defect, which in turn makes me more likely to select the strategy that leaves me least vulnerable to that defection. In other words, *we are very likely to end up in the non-cooperative equilibrium!*

This is a very pessimistic result: we both prefer the cooperative equilibrium to everything else, and this fact is common knowledge. And yet, even small amounts of doubt about the trustworthiness of the other player along with desire to protect oneself from being wrong about the other is almost certain to produce the second worst outcome for both us. In the Prisoner's Dilemma, players are tempted to defect from the cooperative outcome because doing so gives them unambiguous benefit. In the Stag Hunt, this is not so: each player is certain to lose if he unilaterally defects from the cooperative outcome. In both cases, however, mutual defection is likely to happen.

The advantage of a SH-like situation over a PD-like situation is that the social dilemma is solvable in principle in the first case but not in the latter. For instance, if we manage to coordinate expectations and attain a level of trust between ourselves, we will cooperate in SH but still will not cooperate in PD. The cooperative outcome can be sustained in equilibrium in SH but not in PD, which implies that one possible solution to cooperation failure in SH is to work on expectations.

As a social dilemma, SH situations can be seen whenever there is a socially suboptimal equilibrium when there exists another that every player would prefer but that involves an action that is rational only if one expects everyone else to choose it. For instance, female circumcision could be thought of as a SH situation. Clearly, mass circumcision is an equilibrium: all young girls get circumcised, and only those that do can find husbands. Any one girl that fails to be circumcised will not find a husband because men in that culture prefer their wives to be circumcised. Since finding a husband is crucial for the well being of the girl and her family, the girl is better off circumcised (or the family is, in which case they force her to do it).

[2] Evolutionary models in which reproduction rates depend on relative success from interactions also select the risk-dominant equilibrium.

It is well-known that circumcision carries significant risks to health, it mutilates the girls, and renders urinating and sexual intercourse extremely painful for life. Without the threat of remaining a spinster, the girl (and the family) is unlikely to do it. This means that the situation in which no girl is circumcised is also an equilibrium. Since nobody is circumcised, the men cannot demand circumcision as a precondition to marriage: if one suitor does, the girl will simply move on, and his next prospective bride will not be circumcised either. Any single girl that does get circumcised can slightly improve her marriage prospects but since she is very likely to get married without it, the health risks and long-term pain outweigh the marginal gain in that probability. Therefore, no girl (or family) wants to get circumcised.

This, then, is a Stag Hunt scenario. There are two possible equilibria: one in which everybody practices circumcision, and another in which nobody does. If a society is in the perverse equilibrium, it may be very difficult to move to the other one. The reason for that is that while everybody may realize the benefit of the cooperative non-mutilating equilibrium, not getting circumcised only makes sense if everybody thinks that everybody else is not getting circumcised it either (or, at least the vast majority of girls are not). Failure to get circumcised unilaterally carries huge risks and girls (or their families) are very unlikely to do it, which of course reinforces the expectation that others are doing it, and so on. Breaking out of this perverse equilibrium may be extremely difficult. In particular, convincing everyone that female circumcision is extremely detrimental to health is not very likely to work as long as the society is such that marriage is vital for a family's prosperity and the culture is such that men strictly prefer circumcised women for wives.

3.4 The Arms Race as a Stag Hunt

One possible objection to depicting the Arms Race dilemma as a PD is that it seems to require the actors to be aggressive in the sense that they both prefer to compel the other to capitulate than live with the status quo. However, as we shall see later on in the class, even classic antagonists like the U.S. and the USSR could be said to have become essentially status quo powers as the Cold War dragged on. Neither one of them had any great interest in challenging the status quo except at the margins. We could argue that the Arms Race had ceased to be a PD and had become a SH situation.

To represent this situation, let us keep the price of the weapon at \$5 million, the prize at \$15 million, and increase the value of the status quo to \$12 million. As before, the net benefit of victory is $15 - 5 = 10$, the net loss of defeat is -15, and the arms race is the status quo value minus the costs of arming, or $12 - 5 = 7$. Table 4 shows the payoffs that represent the preference ordering, and illustrates that the preferences form a Stag Hunt situation.

Mutual disarmament would be the cooperative outcome which preserves the sta-

	Status Quo	$\succ$	Victory	$\succ$	Arms Race	$\succ$	Defeat
	12	$>$	10	$>$	7	$>$	-15
Us	(C, C)	$\succ$	(D, C)	$\succ$	(D, D)	$\succ$	(C, D)
Them	(C, C)	$\succ$	(C, D)	$\succ$	(D, D)	$\succ$	(D, C)

Table 4: Payoffs in the Arms Race Game as a SH.

tus quo and avoids the expense of building weapons. If the other side is expected to cooperate, then each player prefers to do so as well. On the other hand, if one fails to arm when the other one does, the disarmed player would be saddled with the worst possible outcome: defeat. Prudential reasoning suggests that the less risky choice is to arm: you would get your second-best choice is the opponent is cooperating and you would end up in an arms race if he defects as well. An arms race, while expensive, is much preferable to defeat. Small amounts of suspicion about the opponent's intent would then make $\langle D, D\rangle$ the more likely outcome. Let us verify this. The product of deviation losses from $\langle C, C\rangle$ is $(12-10)(12-10) = 4$, and the corresponding product for $\langle D, D\rangle$ is $(7-(-15))(7-(-15)) = 484$. Thus, $\langle D, D\rangle$ risk-dominant and we should expect players to coordinate on it.

The logic of the arms race in a SH-like scenario is fundamentally one of mistrust, risk-aversion, and prudential reasoning. The logic of an arms race in a PD-like scenario is one of desire to exploit the other side's cooperative effort combine with a desire to avoid being saddled with the worst possible outcome. In this sense, the Stag Hunt is probably captures the dynamics of fear-induced hostility much better than a Prisoner's Dilemma.

In international politics, one cannot know the intent and motivations of one's opponent (or partner). We cannot peek into the heads of decision-makers to verify that they do not intend to attack us, which is (of course) what they usually claim. Intentions are not only unverifiable, they are volatile. Changing governments, the particular mood of the leader, or many other factors may change the evaluation of the desirability of attack on a moment's notice. This is why states normally do not rely on intentions, they are forced to *infer* intent from *observable* capabilities and behavior.

This is where suspicion comes into play. If I cannot be certain that my opponent has no intention to attack me, I must admit the possibility (however small) that he might do so. Since being defeated is the worst possible scenario for me, prudential reasoning might lead me risk losing the cooperative outcome in favor of securing, at the very least, a costly preservation of the status quo. So I build some weapons to guarantee my security. Unfortunately, my act of increasing my security immediately decreases the security of my opponent. He would reason as follows: "I was almost sure that he did not have hostile intent but now I see him arming. I know he claims it is purely for defense but is that so? Perhaps he intends to catch me unprepared and defeat me? And even if that is not so, he clearly does not trust me enough or else he would not have started arming. I would like to reassure him that I can be trusted but the only way to do so is to remain unarmed, which unfortunately is very risky if he does happen to have aggressive intent. So I better arm just to make sure I will not have to surrender in that eventuality."

My opponent then arms as well, which makes me even less secure. We both have matched each other in armaments, the status quo survives, but we also learned that we cannot trust each other not to arm. Because we cannot observe intent, we can only see the arming decision which could be because the other side is afraid or it could be because the other side is aggressive. Reassurance being too risky, we opt for the prudential choice and continue arming, further increasing the suspicion and hostility. The process feeds on itself and rationalizes the non-cooperative outcome, just as in the original Stag Hunt story. The process, in which small doubts lead to defensive measures which increase the insecurity of the opponent, who reacts with defensive measures of his own, which increases my insecurity and as well as my doubts leading to further defensive measures on my part, is called the *Security Dilemma*, and it is very similar to the Stag Hunt scenario.

Notice that once the suspicion starts, it is in the interest of the players to restore trust and get the cooperative equilibrium. Unfortunately, trust can only be restored if one of the players decides to take the risk and plunge into unilateral disarmament. If his opponent turns out to have a SH preference structure (prefers the status quo without arms to victory), then this gesture would be reciprocated and the players could potentially go to a stable cooperative solution. If, on the other hand, one's opponent turns out to have a PD preference structure, then one risks defeat. If one suspects that the opponent has PD preferences or if one's opponent is so suspicious that he would ignore the gesture, no player would make the necessary first step to achieving cooperation.

What model you think represents the Arms Race problem best depends on what you think the structure of the preferences is. If you think of the Arms Race as a Prisoner's Dilemma, you would not recommend trust-building and risky unilateral actions: the opponent is sure to ignore anything you say and would not reciprocate restraint because exploiting your weakness is preferable to cooperation. If you think of the Arms Race as a Stag Hunt, on the other hand, you would recommend trust-building, and you might even recommend a dramatic unilateral gesture that runs serious risks but that can persuade the opponent of your peaceful intent. (We shall see how precisely this type of gesture by the Soviet Union was the catalyst for ending the Cold War.)

These illustrations underscore the major reason for doing this abstract analysis. Once we learn to recognize the equivalence of different strategic situations, we can apply the insights from a model describing one of them directly to another without even having to build a model to represent it. In this course, our goal is to study a series of games to build our intuition about what types of situations seem to occur that concern national security. Once we begin recognizing the similarities (strategic equivalence) between different situations, we can apply our insights to analyze them without actually having to construct explicit models. We shall see that the abstract games tell us quite a bit how to deal with adversaries as disparate as the Soviets, Saddam, or terrorists!

4 Summary

In sum, if we want to analyze a situation, we first try to come up with a model that captures the essential elements of the interaction, and then we "solve" the model, meaning we find its equilibria. An equilibrium is a prescription (what each player is going to do, and what each player expects its opponent to do) that is consistent with rational and intelligent behavior. Why? Because in equilibrium all players are doing their individual best given their goals. Analyzing a model consists of finding strategies for all players that are mutual best responses and therefore make for an equilibrium. The concepts you should remember are as follows:

- **Best response** strategy: a strategy for one player that is the best this player can do given what his opponent is doing. Best responses are always defined in terms of the opponent's strategy (hence, "response"). There may be multiple best responses: that is, a player may have more than one strategy that is best (which means that they all yield the same expected payoff).

- **Strategy profile**: a set of strategies, one for each player. With two-player games, a strategy profile will have two elements. Each strategy profile defines an expected outcome for the game. That is, it tells us what will happen if the players actually follow the strategies specified by the profile.

- **Nash equilibrium**: a strategy profile in which all strategies are best responses to each other. That is, each player's strategy is a best response to what all other players are doing. Alternatively, Nash equilibrium is a strategy profile such that no player wants to change his strategy given the strategies of his opponents (hence the idea of "equilibrium": no player wants to deviate from it).

  – to show that a strategy profile is *not* a Nash equilibrium, all you have to do is find *one* player who could benefit by changing his strategy (so you can stop the analysis after finding even one such instance);

  – to show that a strategy profile is a Nash equilibrium, you have to show that no player wants to change strategy (so you have to examine each and every player's strategy separately).

- **Equilibrium outcome**: the game outcome that will be realized if the players actually follow the strategies prescribed by the equilibrium strategy profile.

Thus far, we have only seen Nash equilibria in pure strategies. Next time we look at how the concept of best responses can be extended to mixed strategies. Nash's original contribution was to (a) define the solution concept that today bears his name, and (b) show that a very large class of games actually have at least one such equilibrium. This made the solution very attractive because it promised to yield answers for a wide variety of substantive problems. As we shall see later on, the solution does have some important shortcomings, so we shall investigate two other solution concepts that are designed to overcome them: "subgame perfect" equilibrium, and "sequential" equilibrium. These are just variants of Nash equilibrium with some additional requirements beyond the strategies being best responses to each other.