

National Security Strategy: Strategies in Extensive Form Games

Professor Branislav L. Slantchev

January 1, 2014

Overview We know how to represent various situations with the help of games in extensive form. We now begin learning how to analyze them. First, we study a more precise definition of a “strategy” as a complete contingent plan of action. We see how to define the strategies in games of perfect and imperfect information and learn to distinguish between pure and mixed strategies. We then see how to do this in games of incomplete information, where we define three kinds of type-contingent strategies: pooling, separating, and semi-separating.

In the previous lecture, we learned how to use game trees to describe formally and abstractly any situation where outcomes depend on the strategic interaction of several players. Obviously, any such description is a stylization at best, it omits a lot of real world specifics in order to make its analysis tractable. We aim to capture the “essential” features of a situation by removing all its “inessential” components.

Clearly, there is no cookie-cutter method of doing this. It’s an art form and it takes a lot of practice, patience, intuition, and insight just to know how to distill a complex real-world interaction into a form we can analyze. Sometimes, analysts miss important features, but when they get nonsensical results from the analysis, they usually become aware of the problem, go back to the specification of the game, and redo it. It’s a long and iterative process, and we shall see how one might go about tackling it.

It is important to remember the concepts of perfect and complete information. Also, you should review the ways of converting games of incomplete information into games of imperfect information. We now begin studying the methods of analysis of strategic situations.

So now we can describe, at least in principle, just about any strategic situation that we might be interested in. What do we do next? We “analyze” the game. Fine, but a bit too vague. Recall what we are ultimately interested in: we want to understand how rational and intelligent actors would behave in such a situation. If we understand that (and the reasons why they do it), then we can design appropriate policies, and interpret historical cases much more fruitfully than we probably could have done without the aid of this tool.

For example, consider the run-up to the last war in Iraq. We knew Saddam Hussein probably preferred to live to being deposed and dying. On the other hand, he also must have preferred standing firm to US demands because revealing the weakness of his regime could have compromised it fatally, because there was a chance the American allies would be able to restrain Bush, and Iraq could still emerge as the sole power that would dare to defy (successfully) the Americans in the region. He must have known that resisting the US would be a dangerous and risky tactic: should the US invade, it was highly likely that his regime would fall. Hence, when he had to decide what would happen if he resisted the US. That, of course, would depend on how the US was expected to react if he did. If it were sure to invade, then resistance would be a bad idea. But if it chose to hew the line with its allies and settle for inspections, then resistance would be a good idea. However, it was not clear what the US would do. War, no matter how victorious, would be quite costly, and hence the US would prefer to settle it peacefully if possible. So when it decides how bellicose to be, the US must consider the likely consequences: if it presses too hard, and Saddam resists, it may have to go to war. If it presses too softly, Saddam would probably resist, and the US would end up without concessions by Iraq. Hence, US strategy must be predicated on expected behavior by Saddam. However, as we just saw above, Saddam’s own behavior is predicated on what the US is expected to do,

which means that US behavior is predicated on what it expects Saddam to expect the US to do, and so on.

In other words, when we say we want to analyze what actors would do in such a situation, we mean that we are interested in finding out what **strategies** they are going to formulate based on these interactive expectations. What is the best course of action each actor would pursue given that its opponent is pursuing its own best course of action? Before we define what it means for a strategy to be the *best course of action*, we have to understand what a strategy is. As we shall see, although much of its definition is intuitive, there are still important differences from the way the word is used in everyday language, so pay attention.

1 Strategies

A strategy is a complete contingent plan of action. It specifies the actions the player is to perform in each possible contingency that might arise in the course of the game. As you should recall, players get to move every time one of their information sets is reached. Therefore, *a strategy must prescribe what action to take at each information set.* Consider Figure 1.




Figure 1: Two Actions Per Player, Perfect and Imperfect Information.

In both cases, we have two players, and each player has two actions. Consider the game of perfect information. The US has one information set (the initial node). According to our definition of a strategy, the strategy for the US would have only one component: what action should the US choose at its single information set. Because there are two possible actions at this information set, the US has two strategies: choose E or choose $\sim E$. We shall write the set of strategies as follows:

$$S_{\text{US}} = \{E, \sim E\}.$$

Note that this is an *unordered list*, which is why we use the curly braces, $\{\dots\}$, to enclose its contents. We can list the strategies in any order we wish, it is not significant.

Turning now to the USSR, note that it has two information sets that are singletons: one follows action E by the US, and the other follows action $\sim E$ by the US. This

means that a strategy for the USSR would have two components, an action at the first information set, and another action at the second information set. For example, one strategy would prescribe $\sim e$ if the US plays E , and e if the US plays $\sim E$. We would write this strategy as the pair $(\sim e, e)$, with the convention that the actions are listed in the order in which information sets appear in the game tree, top to bottom and left to right. That is, the strategy $(e, \sim e)$ prescribes e after the US plays E and $\sim e$ after the US plays $\sim E$.

Note now that the order of actions listed in the strategy does matter, so we have an *ordered list*. That is why we use parentheses, $(. . .)$, to denote its contents. The idea is that order matters because we want to know what action the strategy prescribes for a particular information set. Contrast this with the use of curly braces to list sets of strategies. Here we use parentheses to list the actions specified by one particular strategy.

The USSR has two information sets, with two actions at each set, and so it has a total of $2 \times 2 = 4$ different strategies:

$$S_{\text{USSR}} = \{(e, e), (e, \sim e), (\sim e, e), (\sim e, \sim e)\}.$$

What do these strategies mean? Take, for example, $(\sim e, e)$. It reads “back down if the US escalates and escalate if the US backs down.” Similarly, the strategy $(\sim e, \sim e)$ reads “back down no matter what the US does.” Notice that each of these is a complete contingent plan of action: each strategy specifies what the USSR should do following the move by the US (so one contingency is the US escalating, and the other the US backing down); and each strategy is also complete because it specifies what the USSR should do in every possible contingency in this game. Since this game is one of perfect information (the US move is observable), the USSR can *condition* its behavior on that of the US.

Let’s now take a look at the game of imperfect information. As before, the US has one information set, so its set of strategies is the same. However, now the USSR also has one information set. Accordingly, its strategy has only one component: what to do at this information set. This is important, so remember it. It does not matter how many nodes the information set contains, there is only one action that the player can choose for the information set from the set of available actions there. In this case, the USSR has two actions, so its strategy is simply:

$$S_{\text{USSR}} = \{e, \sim e\}.$$

The strategy space for the USSR now looks very different. Why? Since this is a game of imperfect information, the USSR does not observe the move by the US, and therefore cannot condition its behavior on that. Simply put, there really is only one contingency in this game and it arises when it is USSR’s turn to move without knowing that the US has done. Hence, each strategy specifies an action for that contingency: either escalate or back down. The USSR cannot condition on something it knows nothing about.

To recap, a strategy must specify one action for each information set for the player. Note that in the perfect information game, the instruction for the USSR that says “escalate if the US backs down” cannot be a strategy because it does not say what to do if the US escalates. However, in the imperfect information game, the USSR does not know what action the US has chosen when making its own choice, so its plan is not contingent on US behavior in the sense that the strategy can only prescribe either “escalate” or “back down”.

1.1 Examples: Games of Complete Information

Let’s look at several examples. In Figure 2, the USSR moves first and chooses whether to issue a threat, T , or keep quiet, Q . If it keeps quiet, the status quo prevails, and nothing happens. If it issues a threat, then the players enter the crisis **subgame**, which is the same as before. The US has only one information set, and so its strategies are $S_{US} = \{E, \sim E\}$.




Figure 2: Initial Move by USSR, then Crisis Subgame with Perfect Information.

The USSR has *three* information sets, and so each of its strategies must specify three actions: what to do initially, and then what to do following the possible reactions of the US. For example, one such strategy would be to “threaten, then back down if the US escalates, and escalate if the US backs down,” or, in symbols, the triple $(T, \sim e, e)$. Another strategy would be “threaten and escalate regardless of what the US does”: (T, e, e) .

The following is crucially important. Because the strategy must specify an action for each information set, it means that a **strategy must specify actions for information sets that may not be reached if the strategy is followed**. Here’s an example: $(Q, \sim e, e)$. This strategy reads “keep quiet, back down if the US escalates, and escalate if the US backs down.” Obviously, if the USSR plays Q , then the crisis subgame is never reached. Nevertheless, the strategy must specify what the USSR would do if that subgame is reached. This appears to be redundant, but it is not. Here’s why.

As we shall see, what USSR chooses as its initial action (T or Q) depends on what it expects to happen if it plays T . That is, what will happen in the crisis subgame. To form expectations about that subgame, the USSR must form expectations about the behavior of the US, which in turn depends on what response the USSR

would have to escalation or backing down. Thus, in order to evaluate the possible course of action by the US, the USSR would have to take into account what it itself would do following actions by the US. Thus, even if the USSR chooses Q initially, this choice may be only optimal because of what it expects to happen if it chooses T , and so the strategy must also specify these actions. *A strategy is not just a plan of action for a particular player, it also summarizes what its opponent would expect the player to do.*

In other words, a formal description of a strategy reflects what the opponent thinks one's plan might be. In that regard, it is more than a simple plan of action. As I told you before, we find optimal solutions to these games by analyzing the strategies of all players simultaneously because each optimal strategy depends on the optimal strategy of the opponent. This means that in order to formulate an optimal strategy, one must consider what the opponent's optimal strategy is, which in turn depends on one's optimal strategy. This circular definition means that specifying a strategy reflects (a) how a player is going to behave, and (b) how its opponent expects it to behave.

Back to our example, the USSR has three information sets, with two actions at each information set. Therefore, each strategy would have three components, and there are $2 \times 2 \times 2 = 8$ different strategies:

$$S_{\text{USSR}} = \left\{ (T, e, e), (T, e, \sim e), (T, \sim e, e), (T, \sim e, \sim e), \right. \\ \left. (Q, e, e), (Q, e, \sim e), (Q, \sim e, e), (Q, \sim e, \sim e) \right\}.$$

Suppose we want the crisis subgame to be one of simultaneous moves. Then we have the situation in Figure 3.




Figure 3: Initial Move by USSR, then Crisis Subgame with Imperfect Information.

Strategy set for US is the same because it only has one information set. The USSR has two information sets, one of them is a singleton, and the other contains two nodes. Because there are two information sets, the strategy should prescribe two actions, one for each of them:

$$S_{\text{USSR}} = \left\{ (T, e), (T, \sim e), (Q, e), (Q, \sim e) \right\}.$$

Again, note that the strategy must be a complete plan of action and specify what to do at information sets that may not be reached if the strategy is followed.

2 Mixed Strategies

The strategies we dealt with in the previous section are called **pure strategies** because they specify with certainty what action the player will take at each information set. That is, in a pure strategy, the plan prescribes choosing a particular single alternative from the set of actions at each information set. However, as I noted above, strategies also reflect what one's opponent expects one to play. In many cases, it is not "good" for the player for its opponent to know with certainty what to expect.

For example, in a crisis it may be to one's advantage to keep the opponent guessing about one's next step. Why? Because if the opponent knows that the next move is "back down" then it will probably choose to press its demands, and will therefore get them. However, if it is not certain about one's next move, then its behavior may be different. How do we represent this sort of uncertainty?

We modify our concept of strategy to include not just the pure actions available at each information set, but also probability distributions over these actions. In our crisis example, instead of the strategy specifying "escalate" or "back down" at one's information set, it specifies instead "escalate with probability q " and "back down with probability $1 - q$." Thus, players are allowed to randomize over the alternatives from which they must choose. This is called a **mixed strategy**.

Thinking about strategies as summaries of the opponent's expectations can also help understand what a mixed strategy is: it can be said to reflect the opponent's uncertainty about the player's behavior. In other words, whereas the player himself knows which pure strategy he will play, the opponent is not quite sure—hence the probabilities attached to the various pure strategies. (In fact, Harsanyi has shown that mixed strategies can be given exactly such an interpretation.)

One possible mixed strategy is to escalate with probability $1/3$ and back down with probability $1 - 1/3 = 2/3$. Another is to escalate with probability $.25$ and back down with probability $1 - .25 = .75$. Clearly, the number of possibilities is infinite. It is worth noting that pure strategies are limiting cases of the mixed strategies. For example, the mixed strategy "escalate with probability 1" is equivalent to the pure strategy "escalate," while the mixed strategy "escalate with probability 0" is equivalent to the pure strategy "back down."

Going back to our examples in Figure 4. Consider (b) first. The US has one information set with two actions, and so a mixed strategy would specify the probabilities of each: "escalate with probability q and back down with probability $1 - q$." Because there are only two actions, and because the US must choose one of them, the probabilities must sum to 1.

Similarly, in this imperfect information game, the USSR has one information set with two actions, and so a mixed strategy would simply specify the probabilities of each action at this set: "escalate with probability p and back down with probability $1 - p$." Again, the probabilities must sum to 1 because the USSR must take one of




Figure 4: Mixed Strategies with Two Actions Per Player, Perfect and Imperfect Information.

the two actions.

The situation is a bit more complicated in (a), the game of perfect information, because the USSR has two information sets. Thus, it can randomize at two places, either following escalation by the US or following backing down. So, a mixed strategy must specify two probability distributions, one for each information set. For example, “escalate with probability p if the US escalates and escalate with probability r if the US backs down.” Note that the mixing probabilities are independent between information sets. We would write a mixed strategy for the USSR like this:

$$\sigma_{\text{USSR}} = (p, r).$$

That is, p is the probability of escalating at the first information set, and r is the probability of escalating at the second. Note that because there are only two actions at each set, this specification implicitly gives the probability of backing down at the first information set ($1 - p$), and the second information set ($1 - r$).

Note that the mixed strategy $(0, 1)$ is actually the pure strategy $\{e, \sim e\}$ because it specifies: “escalate with $p = 0$ (or back down) if US escalates, and escalate with $r = 1$ if the US backs down.” Similarly, the mixed strategies $(0, 0)$, $(1, 0)$, and $(1, 1)$ all represent pure strategies.

Thus, a pure strategy specifies what action to take at each information set with certainty. A mixed strategy specifies with what probability to choose an action at each information set. In other words, a mixed strategy specifies one **probability distribution** for each information set of the player. What is a probability distribution? It is a set of probabilities, one number for each possible action at the information set, such that the numbers sum up to one.

2.1 Mixed Strategies Specify Probability Distributions

Let’s use an example to illustrate these ideas. Consider the game in Figure 5. Player 2 has two information sets, one is reached if player 1 chooses L , and the other is reached if player 1 chooses either C or R . This is a game of imperfect information. Note that when an information set is not a singleton (that is, it contains more than

one node), then *the actions emanating from each node in this set must be the same*. The reason for this is intuitive: If one of the nodes had a different number of actions (or different actions) emanating from it, this means that a player could tell which node he is at by looking at which actions are available to him. In other words, the nodes cannot be in the same information set because an information set summarizes the idea that the player does not know at which node he is at.

For example, suppose player 2 could play an additional action A at the node following C (which is contained in the same information set as the node following D). What happens then if this information set is reached (player 1 has chosen either C or R)? Player 2 would look at the options available to him, and if A is among them, he concludes that he is at the node following C . Otherwise, he concludes that he is at the node following R . Because he knows which node he is at precisely, these nodes cannot be part of the same information set. Recall that an information set represents the idea that player 2 does not know whether player 1 has chosen C or R .




Figure 5: A Game of Imperfect Information.

Now, what are the pure strategies for player 2? Since it has two information sets, each pure strategy must have two components: what action to take if player 1 chooses L (first info set), and what action to take if player 1 chooses either C or R (second info set). Because there are two actions for the first set and three for the second, the total number of pure strategies will be $2 \times 3 = 6$. Make sure that you can list them all.¹ The idea is that a pure strategy *lists precisely which action to take at each information set*. In other words, a player is certain to take the action specified by the strategy.

Mixed strategies are different. Instead of specifying a particular action for each information set, they specify a probability distribution over the actions available at this set. A mixed strategy would specify as many probability distributions as there are information sets, so they will have the same number of components as pure strategies. The difference is that a component of a pure strategy is an action, and a component of a mixed strategy is a probability distribution.

¹Here they are: $\{(a, U), (a, M), (a, D), (b, U), (b, M), (b, D)\}$.

Let p be a probability distribution for the first information set (reached after L), and let q be a probability distribution for the second information set (reached after either C or R). So, p must assign probabilities to the actions a and b , while q must assign probabilities to the actions U , M , and D . Note that p and q are unrelated: The player is free to choose different mixtures at different information sets.

One possible mixture at the first set would be ($p(a) = 1/3, p(b) = 2/3$). That is, the probability of playing a , denoted by $p(a)$ is $1/3$. The probability of playing b , denoted by $p(b)$ is $2/3$. The two probabilities must sum up to 1. Note that any numbers here will do as long as every $0 \leq p(a) \leq 1$, and $0 \leq p(b) \leq 1$, and as long as $p(a) + p(b) = 1$. That is, as long as the individual probabilities assigned to actions are all valid probabilities (numbers between zero and one), and as long as they sum up to one. The reason they have to sum up to one is that when the player reaches the information set, he must choose one of the available actions. In other words, some action will be chosen with probability one. Suppose, for example, that $p(a) = p(b) = 1/3$, which means that the probability that player 2 would choose either a or b at its information set is $1/3 + 1/3 = 2/3$. Since there are no other actions available, this implies that with probability $1 - 2/3 = 1/3$, player 2 would take no action whatsoever. But according to our specification of the game, this is not possible. In other words, player 2 must choose something, so the probabilities assigned to individual actions must sum to one.

An example probability distribution q could be:

$$\left(q(U) = 1/2, q(M) = 1/3, q(D) = 1/6 \right) \equiv (1/2, 1/3, 1/6).$$

The expression on the right of this equivalence is just a short-hand way of writing the probability distribution. In fact any numbers will do here as well as long as:

- $0 \leq q(U) \leq 1$, and $0 \leq q(M) \leq 1$, and $0 \leq q(D) \leq 1$, and
- $q(U) + q(M) + q(D) = 1$.

Clearly, there is an infinite number of possible mixtures for each information set. Of particular interest are the *degenerate mixtures*, that is, probability distributions that assign probability 1 to one of the outcomes. For example, the mixture $q = (1, 0, 0)$ says “play U with probability 1, and play M and D with probability zero.” In other words, this is equivalent to what a pure strategy would specify: “play U .” Thus, we can use mixed strategy notation to represent pure strategies.

How do we write a mixed strategy for player 2 in the game in Figure 5? Like the pure strategies, a mixed strategy would have two components: it would specify a probability distribution for the first information set, and another probability distribution for the second information set. Since we have labeled these already, the mixed strategy would be written as: (p, q) , where p and q are probability distributions.

Suppose we want to use $p = (1/4, 3/4)$, that is play a with probability $1/4$, and play b with probability $3/4$. Suppose also we want to use $q = (1/2, 0, 1/2)$. Then, the mixed strategy that specifies these two randomizations would be written as:

$$(p, q) = ((1/4, 3/4), (1/2, 0, 1/2)).$$

Suppose we want to write the mixed strategy: “choose a and b with equal probability if player 1 chooses L , and choose M if player 1 chooses either C or R .” This tells us that $p = (1/2, 1/2)$, and that $q = (0, 1, 0)$, or the complete specification of the mixed strategy:

$$((1/2, 1/2), (0, 1, 0)).$$

Suppose we want to write the pure strategy: “choose b if 1 plays L , and choose D if 1 plays either R or C ” can be written as

$$((0, 1), (0, 0, 1)).$$

Make sure you understand how this is done. In this case, $p = (0, 1)$. That is, $p(a) = 0$ and $p(b) = 1 - p(a) = 1$. Also, $q = (0, 0, 1)$. That is, $q(U) = 0$, $q(M) = 0$, and $q(D) = 1 - q(U) - q(M) = 1$. We then simply list (p, q) .

To summarize, *a pure strategy specifies one action for each information set of the player. Thus, the number of actions specified by a strategy equals the number of information sets for that player. A mixed strategy specifies one probability distribution for each information set of the player. Thus, the number of probability distributions specified by a mixed strategy also equals the number of information sets for that player. When specifying both pure and mixed strategies, all information sets must be included in the list.*

Furthermore, remember that a pure strategy is nothing more than a mixed strategy with degenerate probability distributions for all information sets. Note that a strategy that has at least one non-degenerate probability distribution is mixed, not pure. Thus, in our example above, the mixed strategy $((0, 1), (0, 0, 1))$ is actually the pure strategy (b, D) . But the mixed strategy $((1/2, 1/2), (0, 0, 1))$ is not pure because even though the mixture at the second information set is degenerate (it assigns probability 1 to D), the mixture at the first information set is not (it assigns equal probabilities to a and b).

Mixed strategies reflect the ideas that (a) a player may want to randomize its actions to keep the opponent guessing, and (b) the opponent may not be sure which action the player may select. In the end, one action is always chosen, but until this is done, it is not certain which action it will be.

Mixed strategies are not about a player flipping coins to determine which of the available actions to choose. Rather, they reflect the idea that its opponents may not be quite sure what it is going to do. For example, suppose USSR can be either tough or weak but the US does not know which. The US has a belief that the USSR is tough with probability p , and weak with probability $1 - p$. Suppose

tough opponents always escalate but weak ones always back down. So, each type is playing a pure strategy. However, from the US perspective, the strategy appears mixed: escalation occurs with probability p and backing down with probability $1 - p$. These are simply the probabilities associated with the two types. While each type is playing a pure strategy, the small amount of uncertainty makes the strategy appear mixed to its opponent. You can think about mixed strategies as reflecting such uncertainties.

2.2 Examples: Mixed Strategies in Games of Complete Information

Let's do the other examples. Consider Figure 2. The US has one information set, which is a singleton, so a mixed strategy would just specify the probabilities of E and $\sim E$. The USSR, on the other hand, has three information sets, so its mixed strategy would specify three probability distributions: one randomization for each information set. For example, "play T with probability a ; then play e with probability b if the US plays E ; and play e with probability c if the US plays $\sim E$." Thus, the mixed strategy would be the triple (a, b, c) that specifies these probabilities. For example, the triple $(1, .25, 1/3)$ is the mixed strategy "play T , then play e with probability .25 if the US escalates, and play e with probability $1/3$ if the US backs down." Note again that the mixtures are independent across information sets. As a further example, the mixed strategy $(1, 0, 1)$ is the pure strategy $(T, \sim e, e)$.

In the example in Figure 3, the USSR has two information sets, so its mixed strategies would specify the probabilities of playing T at the first, and playing e at the second information set. For example, $(.5, .5)$ denotes the mixed strategy "play T with probability .5, and then play e with probability .5." Again, the following is the set of pure strategies expressed in our mixed strategy notation:

$$S_{\text{USSR}} = \{(1, 1), (1, 0), (0, 1), (0, 0)\}.$$

Make sure you understand why this is so.

2.3 Pure and Mixed Strategies in a Complicated Game

Let's do another example, where the players may have more than two actions per information set. The USSR proposes to open negotiations over Berlin, and the US can either accept that or reject it. If the USSR does not propose negotiations or if the US rejects them, nothing happens. The USSR then makes one of three proposals, P (peace with East Germany), B (American withdrawal from West Berlin), or W (build a wall). The US can either accept, y , or reject, n any of the proposals.

The USSR has two information sets, with two actions for the first and three actions for the second. Thus, any Soviet pure strategy will have two components, and there will be $2 \times 3 = 6$ different pure strategies:

$$S_{\text{USSR}} = \{(O, P), (O, B), (O, W), (Q, P), (Q, B), (Q, W)\}.$$




Figure 6: The Two-Way Division Game, Perfect Information.

Let a denote the probability of playing O at the first information set. Let b denote the probability of playing P , c the probability of playing B at the second information set. Because within each set the probabilities of all actions must sum to 1, this means that the probability of Q is $1 - a$, and the probability of W is $1 - b - c$. A mixed strategy for the USSR is then the set $(a, (b, c))$, which defines the mixtures for all information sets. Writing out the pure strategies in our mixed strategy notation gives us

$$S_{\text{USSR}} = \{(1, (1, 0)), (1, (0, 1)), (1, (0, 0)), (0, (1, 0)), (0, (0, 1)), (0, (0, 0))\}.$$

Make sure you understand why.

The US has four information sets, with two actions at each set. Therefore, any American pure strategy will have four components, and there will be $2 \times 2 \times 2 \times 2 = 16$ different pure strategies:

$$S_{\text{US}} = \{(A, y, y, y), (A, y, y, n), (A, y, n, n), (A, y, n, y), \\ (A, n, y, y), (A, n, y, n), (A, n, n, n), (A, n, n, y), \\ (R, y, y, y), (R, y, y, n), (R, y, n, n), (R, y, n, y), \\ (R, n, y, y), (R, n, y, n), (R, n, n, n), (R, n, n, y)\}.$$

A mixed strategy for the US must specify four different mixtures. Let d be the probability of playing A at the first set (and so, probability of R is $1 - d$); let e , f , and g denote the probabilities of playing y at each of the remaining sets, from left to right. That is, the US plays y with probability e if the USSR plays P , and plays n with probability $1 - e$, and so on. Thus, a mixed strategy will be denoted by the ordered tuple (d, e, f, g) . Writing the pure strategies in mixed strategy notation actually makes them easier to list:

$$S_{\text{US}} = \{(1, 1, 1, 1), (1, 1, 1, 0), (1, 1, 0, 0), (1, 1, 0, 1), \\ (1, 0, 1, 1), (1, 0, 1, 0), (1, 0, 0, 0), (1, 0, 0, 1), \\ (0, 1, 1, 1), (0, 1, 1, 0), (0, 1, 0, 0), (0, 1, 0, 1), \\ (0, 0, 1, 1), (0, 0, 1, 0), (0, 0, 0, 0), (0, 0, 0, 1)\}.$$

Generally, to calculate the number of pure strategies available to a player, multiply the number of actions at each of its information sets. If the player has three information sets with 2 actions at the first set, 4 actions at the second, and 3 actions at the third, the total number of pure strategies would be $2 \times 4 \times 3 = 24$. Obviously, if there number of possible actions at some information set is infinite, then the number of possible pure strategies for that player is also infinite.

3 Strategies with Incomplete Information

We know how to represent games of incomplete information by converting them to games of imperfect information. In the previous section, we learned how to specify the strategies for the players in games of perfect and imperfect information. Another example can't hurt.

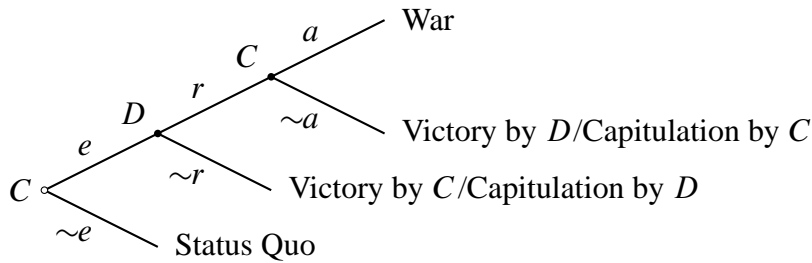


Figure 7: Basic Escalation Game.

Consider the following escalation game. One player, called Challenger (C), begins by choosing to escalate, e , or not, $\sim e$. If it escalates, the Defender (D), can either resist, r , or submit, $\sim r$. If it resists, the challenger can either attack, a , or submit, $\sim a$. The basic sequence is then shown in Figure 7.

Suppose there are two types of challengers: tough and weak. The tough challenger prefers attacking if resisted to submitting, but the weak challenger prefers submitting to attacking. We thus have the preferences, with their numerical representation listed in Table 1.

weak	Victory 10	>	Status Quo 0	>	Capitulation -10	>	War -12
tough	Victory 10	>	Status Quo 0	>	War -1	>	Capitulation -10

Table 1: Preferences of the Challenger against a Weak Defender.

Assume that the defender is weak, so she prefers to capitulate rather than fight. The preferences of the defender are in Table 2. Note that her war payoff depends on the type of challenger she is fighting (she does slightly better against a weak opponent) but in either case, war is worse than capitulation.

if C is weak	Victory 10	>	Status Quo 0	>	Capitulation -10	>	War -12
if C is tough	Victory 10	>	Status Quo 0	>	Capitulation -10	>	War -15

Table 2: Preferences of a Weak Defender.

Using our familiar method of representing this situation with a game tree, we introduce the move by Nature, which determines the type of challenger. The challenger learns its type but the defender does not know it when deciding whether to resist or not. So, Nature chooses the tough challenger with probability p and the weak with probability $1 - p$. The game tree is in Figure 8.




Figure 8: Escalation Game with a Weak Defender and Incomplete Information about the Challenger.

The strategy space for the defender is exceedingly simple. It has only one information set in this game, and therefore its strategy will have only one component specifying what action to take at this set. There are two possible actions, which means there are only two possible pure strategies:

$$S_D = \{r, \sim r\}.$$

The mixed strategy must therefore simply specify the probabilities of playing these actions, say play r with probability q and $\sim r$ with probability $1 - q$.

What are the pure strategies for the challenger? We count information sets. There are two types of challengers, and each type has two information sets. Therefore, a pure strategy for the challenger will have four components: what to do at each of the two information sets if it is tough and what to do at each of the two information sets

if it is weak. For ease of notation, we shall write the strategy as consisting of two pairs of actions, one for each type. For example, the strategy $((e, a), (\sim e, \sim a))$ specifies the following plan of action: “if tough, then escalate and if the defender resists, attack; if weak, then do not escalate and do not attack if the defender resists.” For obvious reasons, this is called a **type-contingent strategy**.

Now, an obvious question would be why we want to specify a type-contingent strategy; that is, why should the tough actor care what its strategy would be if it were weak, after all, the tough actor knows its type, and hence knows that it will never need to know what it would have to do had it been the other type. It would appear that we are being quite redundant here. Or are we? Again, think about the strategies represent: on one hand, they prescribe actions for the player, but on the other hand, they represent what the opponent expects the player to do. As we shall see once we begin analyzing these strategies, the optimality of one’s action at a particular information set depends on how the opponent is expected to react, which in turn depends on one’s own behavior at information sets that follow, and so on. Further note that for each type, the strategy must specify an action for each information set, even ones that are not reached if the strategy as followed. In the example strategy above, the weak type’s plan is to forego escalation (thus ending the game with the status quo) and not attack if it escalates and the defender resists. The logic is equivalent to the one we saw before in games of complete information.

Each of the two types has two information sets with two actions per set, which gives $2 \times 2 = 4$ different pure strategies for each type. With two types and four strategies each, we have a total of $4 \times 4 = 16$ type-contingent pure strategies. The complete set of type-contingent pure strategies for the challenger is:

$$S_C = \left\{ \begin{aligned} &((e, a), (e, a)), ((e, \sim a), (e, \sim a)), ((\sim e, a), (\sim e, a)), ((\sim e, \sim a), (\sim e, \sim a)), \\ &((e, a), (e, \sim a)), ((e, a), (\sim e, a)), ((e, a), (\sim e, \sim a)), \\ &((e, \sim a), (e, a)), ((e, \sim a), (\sim e, a)), ((e, \sim a), (\sim e, \sim a)), \\ &((\sim e, a), (e, \sim a)), ((\sim e, a), (e, a)), ((\sim e, a), (\sim e, \sim a)), \\ &((\sim e, \sim a), (e, \sim a)), ((\sim e, \sim a), (\sim e, a)), ((\sim e, \sim a), (e, a)) \end{aligned} \right\}.$$

These strategies can be divided into two groups: (1) each type plays the same actions at all information sets, or (2) types play different actions at least at one information set.

The first line in the set above specifies the same actions at all information sets regardless of challenger's type. For example, the strategy $((e, a), (e, a))$ prescribes escalation and attack regardless of whether the tough or weak type is concerned. This is called a **pooling strategy** because all types "pool" on the same action. There are four such strategies for the challenger.

Consider now a strategy like $((e, a), (e, \sim a))$. Here, the tough type escalates and attacks, while as the weak type escalates and submits if the defender resists. This is called a **separating strategy** because the two types "separate" themselves by their different courses of action. All the remaining 12 strategies for the challenger in this game are separating.

Intuitively, this is important because if a challenger plays a separating strategy, then the defender may be able to infer the precise type of opponent it is facing from the observable actions of the challenger. Suppose, for example, that the challenger plays $((e, a), (\sim e, \sim a))$. That is, escalate and attack if tough, do not escalate and submit if weak. When the defender gets to move following escalation, it can infer that it is facing the tough challenger for sure. This is because the weak one would not have escalated in the first place. Therefore, the updated belief following escalation will assign $p = 1$ to the probability of the challenger being tough. The challenger **signals** its type (which was private information) by escalating.

Clearly, if both types pool on the same action, it reveals no new information. That is, the defender cannot infer anything about the type of opponent from observing escalation if all types of challenger escalate.

Defining the mixed strategies involves specifying the randomizing probabilities for each information set. For example, let e_t and a_t denote the probabilities with which a tough type escalates and attacks, respectively; and let e_w and a_w denote the corresponding probabilities for the weak type. Then we would write the mixed strategy as $((e_t, a_t), (e_w, a_w))$.

For example, the pooling strategy $((e, a), (e, a))$ can be written in our mixed strategy notation as $((1, 1), (1, 1))$. The separating strategy $((e, a), (\sim e, \sim a))$ can be written as $((1, 1), (0, 0))$.

Consider a mixed strategy like this: "if tough, escalate and attack with probability .9 if resisted; if weak, escalate with probability .5 and back down if resisted." We would write it as $((1, .9), (.5, 0))$. This type of hybrid strategy is called **semi-separating**. This is because the types only partially separate themselves by their actions. For example, after observing escalation, the defender is still unsure whether its opponent is tough because both types escalate with positive probability. (We shall see how the defender would revise its belief with the new information.) However, if the defender observes no escalation, it can conclude that its opponent is weak because only this type fails to escalate with positive probability (the tough type never fails to escalate).

It is worth emphasizing that the strategy $((.8, 1), (.8, 1))$ is pooling, not semi-separating, because the randomizing probabilities are the same for the two types. Thus, no action conveys additional information because all types are equally likely to take it.

As you can see, incomplete information complicates the game quite a bit because it expands the range of possible strategies that players must consider. In this example, the defender has to consider sixteen possibilities for pure strategies and that's only with two types of opponents. In a game of complete information, as in Figure 7, the challenger has only four pure strategies. In all cases, however, the number of mixed strategies is infinite.

In this game, the challenger is the **informed player** because it knows its type. We say that it *possesses private information about its type*. The defender is the **uninformed player** because it does not know something its opponent knows (that's why the information about the type is "private" to the challenger).

When the informed player takes an action, it may reveal some of its private information to the opponent. We say that in this case *signaling occurs*. So in our example, an initial escalation by the challenger may reveal something about its type. To illustrate this, suppose we know that a tough challenger escalates no matter what, but a weak challenger never escalates. Then, if we do observe escalation, we can conclude that the challenger must have been tough because it is the only type that would actually escalate. The defender can *infer* the challenger's type from its action, and *update its belief* that the challenger is tough. Thus, the defender begins the game uncertain about the challenger's type, but can learn something about it from the action the challenger takes. In this case, we say that the challenger *signals*, and that *information is revealed*. Escalation is then a "signal."

When, on the other hand, the uninformed player chooses an action, we say that *screening occurs*. In our example, the defender can resist or submit. Suppose that tough challengers will attack if resisted but weak challengers will back down. The defender can then *screen out* the type of challenger by resisting: If submission occurs, then the defender can conclude that the challenger must have been tough; if escalation occurs, then the defender can conclude that the challenger is weak. Thus, with its action the defender can screen out the different types. Because not resisting involves no further action by the challenger, the defender cannot screen by submitting: It must take an action, whose response by the informed opponent will reveal some of its private information.

This concludes our description of how to specify pure and mixed strategies in various games. Next time we shall learn how to analyze games of perfect information by using the pure strategies and representing such games with tables, called "strategic form," that is easier to analyze.

4 Summary

Strategies are complete contingent plans that specify what action to take at each information set. Thus, they must specify actions at all information sets, including those that are not reached if the player follows the strategy. This necessary because in order to analyze the optimality of an action at a particular information set, one must analyze the consequences of different alternative actions that depend on what the other players are doing, which in turn depends on what this player is planning to do later on.

Pure strategies specify which action to choose at each information set with certainty. Mixed strategies specify probability distributions over these actions for each information set. The probability distributions are independent across information set, that is, a player is free to randomize differently at each of his information sets. A degenerate probability distribution assigns probability 1 to one of the actions. Pure strategies are simply mixed strategies with a degenerate probability distribution for each information set.

An easy way to calculate the number of pure strategies available to a player is to multiply the number of actions at all information sets. It is impossible to calculate the number of mixed strategies: it is always infinity.

Games of incomplete information produce type-contingent strategies for the informed players. That is, they produce a complete plan of action for each possible type. When a player knows its type, it implements the corresponding plan. However, the strategy must specify the plans for all types because the optimality of the behavior of uninformed players will depend on what they think different types are going to do.

When all types of a player have the same plan (even if it includes randomizations), the strategy is pooling. If it differs in at least one component, then it is separating if the different actions are taken with certainty, and semi-separating if at least one type is randomizing. Playing separating and semi-separating strategies reveals information to the uninformed player because signaling can occur. Playing pooling strategies conveys no new information.

When the informed player takes an action that reveals some of its private information, signaling occurs. When the uninformed player takes an action such that its informed opponent's response reveals some of its private information, screening occurs. Signaling and screening are the two of the basic ways players can learn about each other by playing the game.